arXiv:1612.01907v1 [stat.CO] 6 Dec 2016

KFAS: Exponential Family State Space Models in R

Jouni Helske
University of Jyväskylä

Abstract

State space modelling is an efficient and flexible method for statistical inference of a broad class of time series and other data. This paper describes an R package **KFAS** for state space modelling with the observations from an exponential family, namely Gaussian, Poisson, binomial, negative binomial and gamma distributions. After introducing the basic theory behind Gaussian and non-Gaussian state space models, an illustrative example of Poisson time series forecasting is provided. Finally, a comparison to alternative R packages suitable for non-Gaussian time series modelling is presented.

Keywords: R, exponential family, state space models, time series, forecasting, dynamic linear models.

1. Introduction

State space models offer a unified framework for modelling several types of time series and other data. Structural time series, autoregressive integrated moving average (ARIMA) models, simple regression, generalized linear mixed models, and cubic spline smoothing are just some examples of the statistical models which can be represented as a state space model. One of the simplest classes of state space models are linear Gaussian state space models (also known as dynamic linear models), which are analytically tractable, and are therefore often used in many fields of science.

Petris and Petrone (2011) and Tusell (2011) introduce and review some of the contributed R (R Core Team 2016) packages available at Comprehensive R Archive Network (CRAN) for Gaussian state space modelling. Since then, several new additions have emerged in CRAN. Most of these packages use one package or multiple packages reviewed in Tusell (2011) for filtering and smoothing and add new user interfaces and functionality for certain type of models. For example, package **rucm** (Chowdhury 2015) is focused on structural time series, **dlmodeler** (Szymanski 2014) provides a unified interface compatible with multiple packages, and **MARSS** (Holmes, Ward, and Wills 2012, 2013) provides functions for the maximum likelihood estimation of a large class of Gaussian state space models via the EM-algorithm.

One of the packages reviewed in the aforementioned papers is **KFAS** (the Kalman Filtering And Smoothing). Besides of modelling the general linear Gaussian state space models, **KFAS** can also be used in cases where the observations are from other exponential family models, namely binomial, Poisson, negative binomial, and Gamma models.

After the papers by Petris and Petrone (2011) and Tusell (2011), **KFAS** has been completely rewritten. The package is now much more user-friendly due to the use of R's symbolic formulas

in model definition. The non-Gaussian modelling, which was somewhat experimental in the old versions of **KFAS**, is now fully functional supporting multivariate models with different distributions. Many other features have also been added (such as methods for computing model residuals and predictions), the performance of the main functions has improved and in the process several bugs have been fixed.

In this paper I first introduce the basic theory related to state space modelling, and then proceed to show the main aspects of **KFAS** in more detail, illustrate its functionality by applying it to real life dataset, and finally make a short comparison between **KFAS** and other potentially useful packages for non-Gaussian time series modelling.

2. Gaussian state space model

In this section an introduction to key concepts regarding the theory of Gaussian state space modelling as in **KFAS** is given. As the algorithms behind **KFAS** are mostly based on Durbin and Koopman (2012) and the related articles by the same authors, the basic notation is nearly identical with the one used by Durbin and Koopman.

For the linear Gaussian state space model with continuous states and discrete time intervals $t = 1, \ldots, n$, we have

$$\begin{aligned} y_t &= Z_t\alpha_t + \epsilon_t, \quad \text{(observation equation)} \\ \alpha_{t+1} &= T_t\alpha_t + R_t\eta_t, \quad \text{(state equation)} \end{aligned} \tag{1}$$

where $\epsilon_t \sim N(0, H_t), \eta_t \sim N(0, Q_t)$ and $\alpha_1 \sim N(a_1, P_1)$ independently of each other. We assume that y_t is a $p \times 1$, α_{t+1} is an $m \times 1$ and η_t is a $k \times 1$ vector. We also denote $\alpha = (\alpha_1^\top, \ldots, \alpha_n^\top)^\top$ and similarly $y = (y_1^\top, \ldots, y_n^\top)^\top$.

Here y_t contains the observations at time t, whereas α_t is a vector of latent state process at time point t. The system matrices Z_t, T_t, and R_t, together with the covariance matrices H_t and Q_t depend on the particular model definition, and are often time invariant, i.e., do not depend on t. Usually at least some of these matrices contain unknown parameters which need to be estimated. In **KFAS** one defines the model with the function `SSModel`. The function `SSModel` only builds the model and does not perform estimation of unknown parameters, which differs from functions like `lm`, which builds and estimates the model with one command.

The main goal of the state space modelling is to gain knowledge of the latent states α given the observations y. This is achieved by using two important recursive algorithms, the Kalman filtering and smoothing. From the Kalman filtering algorithm we obtain the one-step-ahead predictions and the prediction errors

$$\begin{aligned} a_{t+1} &= \mathsf{E}(\alpha_{t+1}|y_t, \ldots, y_1), \\ v_t &= y_t - Z_t a_t \end{aligned}$$

and the related covariance matrices

$$\begin{aligned} P_{t+1} &= \mathsf{VAR}(\alpha_{t+1}|y_t, \ldots, y_1), \\ F_t &= \mathsf{VAR}(v_t) = Z_t P_t Z_t^\top + H_t. \end{aligned}$$

Using the results of the Kalman filtering, we establish the state smoothing equations running

backwards in time and yielding

$$\hat{\alpha}_t = \mathsf{E}(\alpha_t|y_n, \ldots, y_1),$$
$$V_t = \mathsf{VAR}(\alpha_t|y_n, \ldots, y_1).$$

Similar smoothed estimates can also be computed for the disturbance terms ϵ_t and η_t, and straightforwardly for the signal $\theta_t = Z_t\alpha_t$. For details on these algorithms, see Appendix A and Durbin and Koopman (2012).

A prior distribution of the initial state vector α_1 can be defined as a multivariate Gaussian distribution with mean a_1 and covariance matrix P_1. For an uninformative diffuse prior, one typically sets $P_1 = \kappa\mathrm{I}$, where κ is 10^7, for example. However, this method can be numerically unstable due to cumulative roundoff errors. To solve this issue Koopman and Durbin (2003) present the exact diffuse initialization method, where the diffuse elements in a_1 are set to zero and P_1 is decomposed as $\kappa P_{\infty,1}+P_{*,1}$, where $\kappa \to \infty$. Here $P_{\infty,1}$ is a diagonal matrix with ones on those diagonal elements which relate to the diffuse elements of α_1, and $P_{*,1}$ contains the covariances of the nondiffuse elements of α_1 (and zeros elsewhere). At the start of the Kalman filtering (and at the end of backward smoothing) we use so-called exact diffuse initialisation formulas until $P_{\infty,t}$ becomes a zero matrix, and then continue with the usual Kalman filtering equations. This exact method should be less prone to numerical errors, although they can still occur especially in the smoothing phase, if we, for example, have high collinearity between the explanatory variables of the model. Note that given all the parameters in the system matrices, results from the Kalman filter and smoother are equivalent with Bayesian analysis given the same prior distribution for α_1.

When we have multivariate observations, it is possible that in the diffuse phase the matrix F_t is not invertible, and the computation of a_{t+1} and P_{t+1} becomes impossible. On the other hand, even if F_t is invertible, the computations can become slow when the dimensionality of F_t, that is, the number of series increases. Also in the case of multivariate observations, the formulas relating to the diffuse initialization become cumbersome. Based on the ideas of Anderson and Moore (1979), a complete univariate approach for filtering and smoothing was introduced by Koopman and Durbin (2000) (known as sequential processing by Anderson and Moore). The univariate approach is based on the alternative representation of the model (1), namely

$$\begin{aligned}
y_{t,i} &= Z_{t,i}\alpha_{t,i} + \epsilon_{t,i}, \quad i = 1, \ldots, p_t, \quad t = 1, \ldots, n,\\
\alpha_{t,i+1} &= \alpha_{t,i}, \quad i = 1, \ldots, p_t - 1,\\
\alpha_{t+1,1} &= T_t\alpha_{t,p_t} + R_t\eta_t, \quad t = 1, \ldots, n,
\end{aligned}$$

and $a_{1,1} \sim N(a_1, P_1)$, with the assumption that H_t is diagonal for all t. Here the dimension of the observation vector y_t can vary over time and therefore missing observations are handled straightforwardly by adjusting the dimensionality of y_t. In the case of non-diagonal H_t, the original model can be transformed either by taking the LDL decomposition of H_t, and multiplying the observation equation with the L_t^{-1}, so $\epsilon_t^* \sim N(0, D_t)$, or by augmenting the state vector with ϵ, when Q_t becomes block diagonal with blocks Q_t and H_t. Augmenting can also be used for introducing a correlation between ϵ and η. Both the LDL decomposition and the state vector augmentation are supported in **KFAS**.

In theory, when using the univariate approach, the computational costs of filtering and smoothing decrease, as the number of matrix multiplications decrease, and there is no need for solving the system of equations (Durbin and Koopman 2012, p. 159). As noted in Tusell

(2011), these gains can somewhat cancel out as more calls to linear algebra functions are needed and the memory management might not be as effective as working with larger objects at once. Nevertheless, as noted previously, sequential processing has also other clear benefits, especially with diffuse initialization where the univariate approach simplifies the recursions considerably (Durbin and Koopman 2012).

KFAS uses this univariate approach in all cases. Although v_t, F_t, and $K_t = P_t Z_t^\top = \mathsf{COV}(a_t, y_t | y_{t-1}, \ldots, y_1)$ differ from the standard multivariate versions, we get $a_t = a_{t,1}$ and $P_t = P_{t,1}$ by using the univariate approach. If standard multivariate matrices F_t and K_t are needed for inference, they can be computed later from the results of the univariate filter. As the covariances $F_{*,i,t}$, $K_{*,i,t}$, and $P_{*,t}$ relating to the diffuse phase (see Appendix A) coincide with the nondiffuse counterparts if $F_{\infty,i,t} = 0$, the asterisk is dropped from the variable names in **KFAS**, and, for example, the variable `F` is an $n \times p$ array containing $F_{*,i,t}$ and $F_{i,t}$, whereas `Finf` is an $n \times d$, where d is the last time point before the diffuse phase ended.

2.1. Log-likelihood of the Gaussian state space model

The Kalman filter equations can be used for computing the log-likelihood, which in its standard form is

$$\log L = -\frac{np}{2} \log 2\pi - \frac{1}{2} \sum_{t=1}^{n} (\log |F_t| + v_t' F_t^{-1} v_t).$$

In the case of the univariate treatment and diffuse initialization, the diffuse log-likelihood can be written as

$$\log L_d = -\frac{1}{2} \sum_{t=1}^{n} \sum_{i=1}^{p_t} w_{i,t},$$

where

$$w_{i,t} = \begin{cases} \log F_{\infty,i,t}, & \text{if } F_{\infty,i,t} > 0, \\ I(F_{i,t} > 0)(\log 2\pi + \log F_{i,t} + v_{i,t}^2 F_{i,t}^{-1}), & \text{if } F_{\infty,i,t} = 0. \end{cases}$$

See Durbin and Koopman (2012, Chapter 7) for details. Francke, Koopman, and De Vos (2010) show that there are cases where the above definition of diffuse log-likelihood is not optimal. Without going into the details, if system matrices Z_t or T_t contain unknown parameters in their diffuse parts, the diffuse likelihood is missing one term which depends on those unknown parameters. Francke *et al.* (2010, p.411–412) present a recursive formula for computing this extra term, which is also supported by **KFAS**.

2.2. Example of Gaussian state space model

Now the theory of previous sections is illustrated via example. Our time series consists of yearly alcohol-related deaths per 100,000 persons in Finland for the years 1969–2007 in the age group of 40–49 years (Figure 1). The data is taken from Statistics Finland (2014a,b).

For the observations $y_1, \ldots, y_n$ we assume that $y_t \sim N(\mu_t, \sigma_\epsilon)$ for all $t = 1, \ldots, n$, where μ_t is a random walk with drift process

$$\mu_{t+1} = \mu_t + \nu + \eta_t$$

with $\eta_t \sim N(0, \sigma_\eta^2)$. Assume that we have no prior information about the initial state μ_1 or

the constant slope ν. This model can be written in a state space form by defining

$$Z = \begin{pmatrix} 1 & 0 \end{pmatrix}, H = \sigma^2_\epsilon, T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix},$$

$$\alpha_t = \begin{pmatrix} \mu_t \\ \nu_t \end{pmatrix}, R = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, Q = \sigma^2_\eta,$$

$$a_1 = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, P_{*,1} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, P_{\infty,1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

In **KFAS**, this model can be written with the following code. For illustrative purposes we define all the system matrices manually without resorting default values.

```
R> data("alcohol")
R> deaths <- window(alcohol[, 2], end = 2007)
R> population <- window(alcohol[, 6], end = 2007)
R> Zt <- matrix(c(1, 0), 1, 2)
R> Ht <- matrix(NA)
R> Tt <- matrix(c(1, 0, 1, 1), 2, 2)
R> Rt <- matrix(c(1, 0), 2, 1)
R> Qt <- matrix(NA)
R> a1 <- matrix(c(1, 0), 2, 1)
R> P1 <- matrix(0, 2, 2)
R> P1inf <- diag(2)
R>
R> model_gaussian <- SSModel(deaths / population ~ -1 +
+      SSMcustom(Z = Zt, T = Tt, R = Rt, Q = Qt, a1 = a1, P1 = P1,
+        P1inf = P1inf),
+    H = Ht)
```

The first argument to the `SSModel` function is the formula which defines the observations (left side of tilde operator `~`) and the structure of the state equation (right side of tilde). Here `deaths / population` is a univariate time series, and the state equation is defined using the system matrices with auxiliary function `SSMcustom`, and the intercept term is omitted with `-1` in order to keep the model identifiable. The observation level variance is defined via the argument `H`. The `NA` values represent the unknown variance parameters σ^2_ϵ and σ^2_η which can be estimated using the function `fitSSM`. After estimation, the filtering and smoothing recursions are performed using the `KFS` function.

```
R> fit_gaussian <- fitSSM(model_gaussian, inits = c(0, 0), method = "BFGS")
R> out_gaussian <- KFS(fit_gaussian$model)
```

In this case, the maximum likelihood estimates are 9.5 for σ^2_ϵ and 4.3 for σ^2_η.

From the Kalman filter algorithm we get one-step-ahead predictions for the states $a_t = (\mu_t, \nu_t)^\top$. Note that even though the slope term ν was defined as time-invariant ($\nu_t = \nu$) in our model, it is recursively estimated by the Kalman filter. Thus at each time point t when the new observation y_t becomes available, the estimate of ν is updated to take account of the

new information given by y_t. At the end of Kalman filtering, a_{n+1} gives our final estimate of the constant slope term given all of our data. Here the slope term is estimated as 0.84 with standard error 0.34. For μ_t, the Kalman filter gives the one-step-ahead predictions, but as the state is time-varying, we need to run also the smoothing algorithm if we are interested in the estimates of μ_t for $t = 1, \ldots, n$ given all the data.

Figure 1 shows the observations with one-step-ahead predictions (red) and smoothed (blue) estimates of the random walk process μ_t. Notice the typical pattern; at the time t the Kalman filter computes the one-step-ahead prediction error $v_t = y_t - \mu_t$, and uses this and the the previous prediction to correct the prediction for the next time point (see Appendix A for the detailed update formula). Here this is most easily seen at the beginning of the series where our predictions seem to be lagging the observations by one time step. On the other hand, the smoothing algorithm takes account of both the past and the future values at each time point, thus producing more smoothed estimates of the latent process.

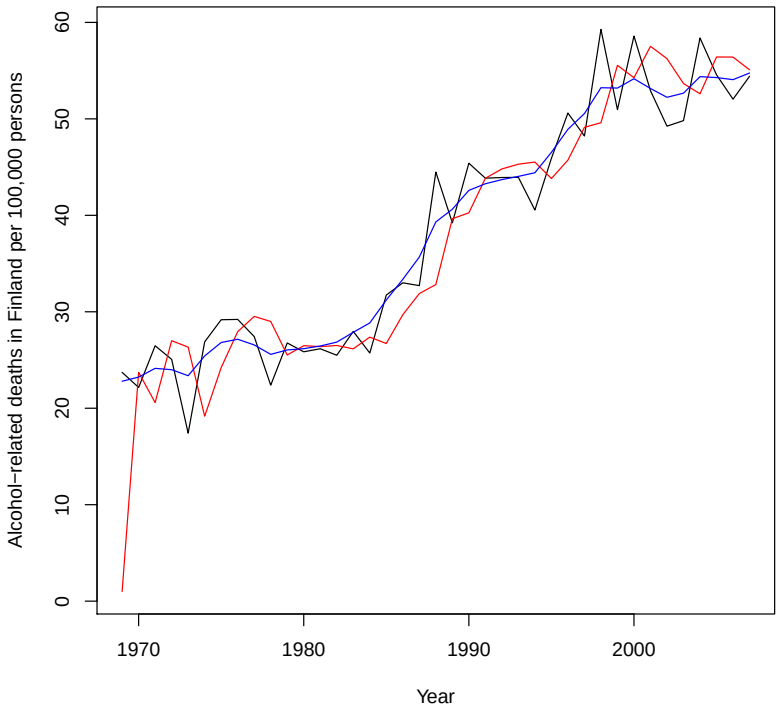

Figure 1: Alcohol-related deaths in Finland in the age group of 40–49 years (black line) with predicted (red) and smoothed (blue) estimates.

3. State space models for the exponential family

KFAS can also deal with observations which come from distributions of an exponential family class other than Gaussian. We assume that the state equation is as in the Gaussian case, but

the observation equation has the form

$$p(y_t|\theta_t) = p(y_t|Z_t\alpha_t),$$

where $\theta_t = Z_t\alpha_t$ is the signal and $p(y_t|\theta_t)$ is the observational density.

The signal θ_t is the linear predictor which is connected to the expected value $E(y_t) = \mu_t$ via a link function $l(\mu_t) = \theta_t$. In **KFAS**, the following distributions and links are available:

1. Gaussian distribution with mean μ_t and variance u_t with identity link $\theta_t = \mu_t$.
2. Poisson distribution with intensity λ_t and exposure u_t together with log-link $\theta_t = \log(\lambda_t)$. Thus we have $\mathsf{E}(y_t|\theta_t) = \mathsf{VAR}(y_t|\theta_t) = u_t e^{\theta_t}$.
3. Binomial distribution with size u_t and probability of success π_t. **KFAS** uses logit-link so $\theta_t = \text{logit}(\pi_t)$ resulting $\mathsf{E}(y_t|\theta_t) = u_t\pi_t$ and $\mathsf{VAR}(y_t|\theta_t) = u_t(\pi_t(1-\pi_t))$.
4. Gamma distribution with a shape parameter u_t and an expected value μ_t, again with log-link $\theta_t = \log(\mu)$, where Gamma distribution is defined as

$$p(y_t|\mu_t, u_t) = \frac{u_t^{u_t}}{\Gamma(u_t)}\mu_t^{-u_t}y_t^{u_t-1}e^{\frac{y_t u_t}{\mu_t}}.$$

   This gives us $\mathsf{E}(y_t|\theta_t) = e^{\theta_t}$ and $\mathsf{VAR}(y_t|\theta_t) = e^{2\theta_t}/u_t$.
5. Negative binomial distribution with a dispersion parameter u_t and an expected value μ_t with log-link $\theta_t = \log(\mu_t)$, where the negative binomial distribution is defined as

$$p(y_t|\mu_t, u_t) = \frac{\Gamma(y_t + u_t)}{\Gamma(u_t)y_t!}\frac{\mu_t^{y_t}u_t^{u_t}}{(\mu_t + u_t)^{u_t+y_t}},$$

   giving us $\mathsf{E}(y_t|\theta_t) = e^{\theta_t}$ and $\mathsf{VAR}(y_t|\theta_t) = e^{\theta_t} + e^{2\theta_t}/u_t$.

Note that the variable u_t has a different meaning depending on the distribution it is linked to. In **KFAS** one defines the distribution for each time series via argument `distribution` and the additional known parameters u_t corresponding to each series as columns of the matrix `u`.

In order to make inferences of the non-Gaussian models, we first find a Gaussian model which has the same conditional posterior mode as $p(\theta|y)$ (Durbin and Koopman 2000). This is done using an iterative process with Laplace approximation of $p(\theta|y)$, where the updated estimates for θ_t are computed via the Kalman filtering and smoothing from the approximating Gaussian model. In the approximating Gaussian model the observation equation is replaced by

$$\tilde{y}_t = Z_t\alpha_t + \epsilon_t, \quad \epsilon_t \sim N(0, H_t)$$

where the pseudo-observations $\tilde{y}_t$ variances H_t are based on the first and second derivatives of $\log p(y_t|\theta_t)$ with respect to θ_t (Durbin and Koopman 2000).

Final estimates $\hat{\theta}_t$ correspond to the mode of $p(\theta|y)$. In the Gaussian case the mode is also the mean. In cases listed in (1)-(5) the difference between the mode and the mean is often negligible. Nevertheless, we are usually more interested in μ_t than in the linear predictor θ_t. As the link function is non-linear, direct transformation $\hat{\mu}_t = l^{-1}(\hat{\theta}_t)$ introduces some bias. To

solve this problem **KFAS** also contains methods based on importance sampling, which allows us to correct these possible approximation errors. With the importance sampling technique we can also compute the log-likelihood and the smoothed estimates for $f(\alpha)$, where f is an arbitrary function of states, $\exp(Z_t\alpha_t)$ being a typical example.

In the importance sampling scheme, we first find the approximating Gaussian model, simulate the states α^i from this Gaussian model and then compute the corresponding weights $w_i = p(y|\alpha^i)/g(y|\alpha^i)$, where $p(y|\alpha^i)$ represents the conditional non-Gaussian density of the original observations, and $g(y|\alpha^i)$ is the conditional Gaussian density of the pseudo-observations $\tilde{y}$. These weights are then used for computing

$$\mathsf{E}(f(\alpha)|y) = \frac{\sum_{i=1}^N f(\alpha^i)w_i}{\sum_{i=1}^N w_i}.$$

The simulation of Gaussian state space models in **KFAS** is based on the simulation smoothing algorithm by Durbin and Koopman (2002). In order to improve simulation efficiency, **KFAS** can use two antithetic variables in the simulation algorithms. See Durbin and Koopman (2012, p. 265-266) for details on how these are constructed.

KFAS also provides means for the filtering of non-Gaussian models. This is achieved by sequentially using the smoothing scheme for $(y_1, \ldots, y_t), t = 1 \ldots, n$ with y_t set as missing. This is a relatively slow procedure for large models, as the importance sampling algorithms need to be performed n times, although the first steps are much faster than the one using the whole data. The non-Gaussian filtering is mainly for the computation of recursive residuals (see Section 4) and for illustrative purposes, where computational efficiency is not that important. With large models or online-filtering problems, one is recommended to use a proper particle filter approach, which is out of the scope of this paper.

For non-Gaussian exponential family models in the context of generalized linear models, a typical way of obtaining the *confidence* interval of the prediction is to compute confidence intervals in the scale of a linear predictor, and then the interval is transformed to the scale of observations. The issue of prediction intervals is often dismissed. For obtaining proper prediction intervals in the case of non-Gaussian state space models, the following algorithm is used in **KFAS**.

(1) Draw N replicates of the linear predictor θ from the approximating Gaussian density $g(\theta|y)$ with importance weights $p(y|\theta)/g(y|\theta)$. Denote this sample $\tilde{\theta}^1, \ldots, \tilde{\theta}^N$ as $\tilde{\theta}$

(2) Using the importance weights as sampling probabilities, draw a sample of size N with replacement from $\tilde{\theta}$. We now have N independent draws from $p(\theta|y)$.

(3) For each $\tilde{\theta}^i$ sampled in step (2), take a random sample of y^i from the observational distribution $p(y|\theta^i)$.

(4) Compute the the prediction intervals as empirical quantiles from $y^1, \ldots, y^N$.

Assuming all the model parameters are known, these intervals coincide (within the Monte Carlo error) with the ones obtained from Bayesian analysis using the same priors for states.

3.1. Log-likelihood of the non-Gaussian state space model

The log-likelihood function for the non-Gaussian model can be written as (Durbin and Koopman 2012, p. 272)

$$\begin{aligned}\log L(y) &= \log \int p(\alpha, y) \mathrm{d}\alpha \\ &= \log L_g(y) + \log E_g \left[\frac{p(y|\theta)}{g(y|\theta)} \right],\end{aligned}$$

where $L_g(y)$ is the log-likelihood of the Gaussian approximating model and the expectation is taken with respect to the Gaussian density $g(\alpha|y)$. The expectation can be approximated by

$$\log E_g \left[\frac{p(y|\theta)}{g(y|\theta)} \right] \approx \log \frac{1}{N} \sum_{i=1}^{N} w_i. \tag{2}$$

In many cases, a good approximation of the log-likelihood can be computed without any simulation, by setting $N = 0$ and using the mode estimate $\hat{\theta}$ from the approximating model.

In practice (2) suffers from the fact that $w_i = p(y|\theta^i)/g(y|\theta^i)$ is numerically unstable; when the number of observations is large, the discrete probability mass function $p(y|\theta^i)$ tends to zero, even when the Gaussian density function $g(y|\alpha^i)$ does not. Therefore it is better to redefine the weights as

$$w_i^* = \frac{p(y|\theta^i)/p(y|\hat{\theta})}{g(y|\theta^i)/g(y|\hat{\theta})}.$$

The log-likelihood is then computed as

$$\log \hat{L}(y) = \log L_g(y) + \log \hat{w} + \log \frac{1}{N} \sum_{i=1}^{N} w_i^*,$$

where $\hat{w} = p(y|\hat{\theta})/g(y|\hat{\theta})$.

3.2. Example of non-Gaussian state space model

The alcohol-related deaths of Section 2.2 can also be modelled naturally as a Poisson process. Now our observations y_t are the actual counts of alcohol-related deaths in year t, whereas the varying population size is taken account of by the exposure term u_t. The state equation remains the same, but the observation equation is now of form $p(y_t|\mu_t) = \text{Poisson}(u_t e^{\mu_t})$.

```
R> model_poisson <- SSModel(deaths ~ -1 +
+      SSMcustom(Z = Zt, T = Tt, R = Rt, Q = Qt, P1inf = P1inf),
+    distribution = "poisson", u = population)
```

Compared to the Gaussian model of Section 2.2, we now need to define the distribution of the observations using the argument `distribution` (which defaults to `"gaussian"`). We also define the exposure term via the argument `u` (for non-Gaussian models the `H` is omitted and vice versa), and use default values for `a1` and `P1` in the `SSMcustom`.

In this model there is only one unknown parameter, σ_η^2. This is estimated as 0.0053, but the actual values of σ_η^2 between the Gaussian and Poisson models are not directly comparable

as the intepretation of μ_t differs between models. The slope term of the Poisson model is estimated as 0.022 with standard error 1.4×10^{-4}, corresponding to the 2.3% yearly increase in deaths.

Figure 2 shows the smoothed estimates of the intensity (deaths per 100,000 persons) modelled as Gaussian process (blue), and as a Poisson process (red).

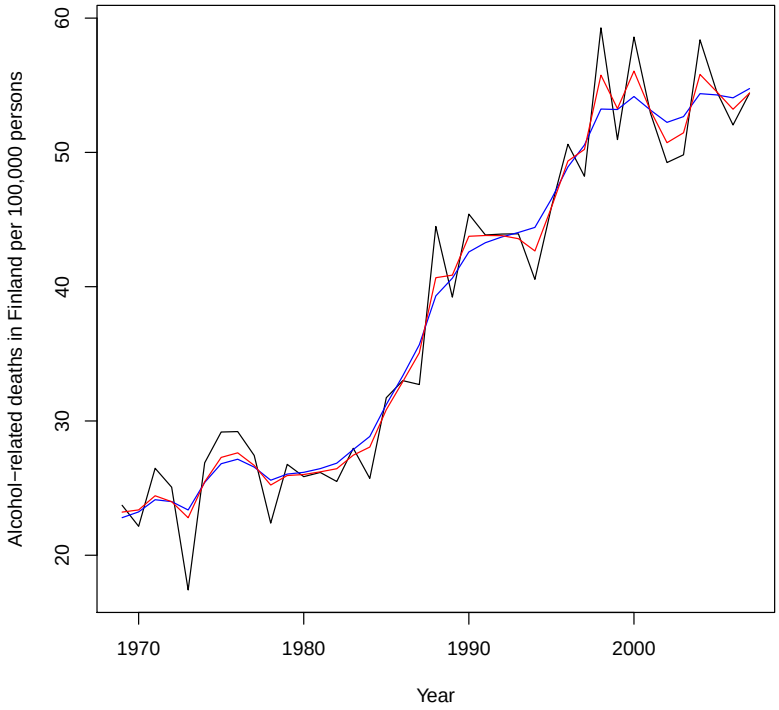

Figure 2: Alcohol-related deaths in Finland (black line) with smoothed estimates from Gaussian model (blue) and Poisson model (red).

4. Residuals

For exponential family state space models, multiple types of residuals can be computed. Probably the most useful ones are standardized recursive residuals, which are based on the one-step-ahead predictions from the Kalman filter. For the univariate case these are defined as

$$\frac{y_t - \mathsf{E}(y_t|y_{t-1},\ldots,y_1)}{\sqrt{\mathsf{VAR}(y_t|y_{t-1},\ldots,y_1)}}, \quad t = d+1\ldots,n,$$

where d is the last time point of the diffuse phase, and the denominator can be decomposed as

$$\begin{aligned}\mathsf{VAR}(y_t|y_{t-1},\ldots,y_1) &= \mathsf{VAR}(\mathsf{E}(y_t|\theta_t,y_{t-1},\ldots,y_1)|y_{t-1},\ldots,y_1)\\ &\quad + \mathsf{E}(\mathsf{VAR}(y_t|\theta_t,y_{t-1},\ldots,y_1)|y_{t-1},\ldots,y_1)\\ &= \mathsf{VAR}(\mathsf{E}(y_t|\theta_t)|y_{t-1},\ldots,y_1) + \mathsf{E}(\mathsf{VAR}(y_t|\theta_t)|y_{t-1},\ldots,y_1).\end{aligned}$$

In the Gaussian case this simplifies to $v_t F_t^{-\frac{1}{2}}$.

For multivariate observations we have several options on how to standardize the residuals. The most common one is a marginal standardization approach, where each residual series is divided by its standard deviation, so we get residual series which should not exhibit any autocorrelations. Another option is to use, for example, Cholesky decomposition for the prediction error covariance matrix F_t and standardize the residuals by $L_t^{-1}(y_t - \hat{y}_t)$ where $L_t L_t^\top = F_t$. Now the whole series of residuals (treated as a single univariate series) should not contain any autocorrelation.

For computing the marginally standardized residuals, multivariate versions of F_t and v_t are needed, whereas the Cholesky standardized residuals can be computed directly from the sequential Kalman filter as

$$v_{i,t} F_{i_t}^{-\frac{1}{2}}, \quad j = 1, \ldots, p, \quad t = d + 1 \ldots, n.$$

These multivariate residuals depend on the ordering of the series, so if the residual diagnostics exhibit deviations from model assumptions, then the interpretation is somewhat more difficult than when using the marginal residuals. Therefore marginal residuals might be preferred. Note that if we want quadratic form residuals $(y_t - \hat{y}_t) F_t^{-1} (y_t - \hat{y}_t)$, then the ordering of the series does not matter.

The recursive residuals are defined just for the non-diffuse phase, which is problematic if the model contains a long diffuse phase, for example, because a dummy variable with a diffuse prior is incorporated to the model. This is because the diffuse phase cannot end before the dummy variable changes its value at least once. In order to circumvent this, one can use a proper but highly non-informative prior distribution for the intervention variable when computing the residuals, which should have a negligible effect on the visual inspection of the residual plots.

Other potentially useful residuals are auxiliary residuals, which are based on smoothed values of states. For details, see Harvey and Koopman (1992) and Durbin and Koopman (2012, Chapter 7).

5. Functionality of KFAS

The state space model used with **KFAS** is built using the function `SSModel`. The function uses R's formula object in a similar way to that of the functions `lm` and `glm`, for example. In order to define the different components of the state space model, auxiliary functions `SSMtrend`, `SSMseasonal`, `SSMcycle`, `SSMarima`, `SSMregression` are provided. These functions can be used to define the structural, ARIMA, and regression components of the model. The function `SSMcustom` can be used for constructing an arbitrary component by directly defining the system matrices of the model (1). More details on how to construct common state space models with **KFAS** are presented in Section 6.

The function `SSModel` returns an object of class `SSModel`, which contains the observations `y` as the `ts` object, system matrices `Z`,`H`,`T`,`R`,`Q` as arrays of appropriate dimensions, together with matrices `a1`, `P1`, and `P1inf` defining the initial state distribution. Additional components contain the system matrix `u` which is used in non-Gaussian models for additional parameters, the character vector `distribution` which defines the distributions of the observations

(multivariate series can have different distributions), and the tolerance parameter `tol` which is used in diffuse phase for checking whether F_∞ is nonzero.

`SSModel` object also contains some attributes, namely, integer valued attributes `p`,`m`,`k`, and `n` which define the dimensions of the system matrices, character vectors `state_types` and `eta_types` which define the elements of α_t and η_t, and integer vector `tv` which defines whether the model contains time-varying system matrices. These attributes are used internally by **KFAS**, although the user can carefully modify them if needed. For example, if the user wishes to redefine the error term η_t by changing the dimensions of `R` and `Q`, the attributes `k` and `eta_types` need to be updated accordingly.

The unknown model parameters can be estimated with `fitSSM`, which is a wrapper around **R**'s `optim` function and the `logLik` method for the **SSModel** object. For `fitSSM`, the user gives the model object, initial values of unknown parameters and a function `updatefn`, which is used to update the model given the parameters (the help page of `fitSSM` gives an example of `updatefn`). As the numerical optimization routines update the model and compute the likelihood thousands of times, the user is encouraged to build his own problem-specific model updating function for maximum efficiency. By default, `fitSSM` estimates the `NA` values in the time invariant covariance matrices H and Q, but no general estimation function is provided. Of course, the user can also directly use the `logLik` method for computing the likelihood and thus is free to choose a suitable optimization method for his problem.

The function `KFS` computes the filtered (one-step-ahead prediction) and smoothed estimates for states, signals, and the values of the inverse link function (expected value μ or probability π) in a non-Gaussian case. For Gaussian models, disturbance smoothing is also available.

With `simulateSSM` the user can simulate the states, signals or disturbances of the Gaussian state space models given the model and the observations. If the model contains missing observations, these can also be simulated by `simulateSSM` in a similar way. It is also possible to simulate states from predictive distributions $p(\alpha_t|y_1,\ldots,y_{t-1})$, $t = 1,\ldots,n$. For these simulations, instead of using marginal distributions $N(a_t, P_t)$, **KFAS** uses a modification of Durbin and Koopman (2002), where smoothing is replaced by filtering.

For non-Gaussian models, `importanceSSM` returns the states or signals simulated from the approximating Gaussian model, and the corresponding weights w_i, which can then be used to compute arbitrary functions of the states or signals.

There are several `S3` methods available for `SSModel` and `KFS` objects. For both objects, simple `print` methods are provided, and for `SSModel` objects there is the `logLik` method. The `predict` method is for computing the point predictions together with confidence or prediction intervals. The extraction operator `[` for extracting and replacing the subsets of model elements is available for the `SSModel` class. Using this method when modifying the model is suggested instead of a common list extractor `$`, as the latter can accidentally modify the dimensions of the corresponding model matrices. A simple `plot` method for residual inspection is also provided.

For the `KFS` object, the methods `residuals`, `rstandard`, and `hatvalues` are provided. Also, a function `signal` can be used for extracting subsets of signals from `KFS` objects, for example, the part of $Z_t\alpha_t$ that corresponds to the regression part of the model.

Methods `coef` and `fitted` for the quick extraction of state or mean estimates are also available for `KFS` and `SSModel` objects.

6. Constructing common state space models with KFAS

This section presents some typical models which can be formulated in a state space form. More examples can be found on the main help page of **KFAS** by typing `?KFAS` after the package is loaded via `"library("KFAS")`. These examples include most of the examples presented in Durbin and Koopman (2012). Additional examples illustrating the functionality of **KFAS** can be found from the documentation of the particular functions.

All the auxiliary functions used in the formula argument of the function `SSModel` have some common arguments which are not directly related to the system matrices of the corresponding component. In complex multivariate models, an important argument is `index`, which defines the series for which the corresponding component is constructed. For example, if we have four time series ($p = 4$), we may want to use a certain regression component only for series 2 and 4. In this case we use the argument `index = c(2,4)` when calling the appropriate `SSMregression` function. By default the index is `1:p` so the component is constructed for all series.

Another argument used in several auxiliary functions is `type`, which can take two possible values. The value `"distinct"` defines the component separately for each series defined by `index` (with covariance structure defined by the argument `Q`), whereas the value `"common"` constructs a single component which applies to all series defined by `index`. For example, we can define distinct random walk components for all series together with a covariance matrix which captures the dependencies of the different series, or we can define just a single random walk component which is common to all series.

6.1. Structural time series

A structural time series refers to the class of state space models where the observed time series is decomposed into several underlying components, such as trend and seasonal effects. The basic structural time series model is of the form

$$\begin{aligned} y_t &= \mu_t + \gamma_t + c_t + \epsilon_t, \quad \epsilon_t \sim N(0, H_t), \\ \mu_{t+1} &= \mu_t + \nu_t + \xi_t, \quad \xi_t \sim N(0, Q_{\text{level},t}), \\ \nu_{t+1} &= \nu_t + \zeta_t, \quad \zeta_t \sim N(0, Q_{\text{slope},t}), \end{aligned} \tag{3}$$

where μ_t is the trend component, γ_t is the seasonal component and c_t is the cycle component. The seasonal component with period s can be defined in a dummy variable form

$$\gamma_{t+1} = -\sum_{j=1}^{s-1} \gamma_{t+1-j} + \omega_t, \quad \omega_t \sim N(0, Q_{\text{seasonal},t}),$$

or in a trigonometric form, where

$$\begin{aligned} \gamma_t &= \sum_{j=1}^{\lfloor s/2 \rfloor} \gamma_{j,t}, \\ \gamma_{j,t+1} &= \gamma_{j,t} \cos \lambda_j + \gamma^*_{j,t} \sin \lambda_j + \omega_{j,t}, \\ \gamma^*_{j,t+1} &= -\gamma_{j,t} \sin \lambda_j + \gamma^*_{j,t} \cos \lambda_j + \omega^*_{j,t}, \quad j = 1, \ldots, \lfloor s/2 \rfloor, \end{aligned}$$

with $\omega_{j,t}$ and $\omega^*_{j,t}$ being independently distributed variables with $N(0, Q_{\text{seasonal},t})$ distribution and $\lambda_j = 2\pi j/s$.

The cycle component with period s is defined as

$$c_{t+1} = c_t \cos \lambda_c + c_t^* \sin \lambda_c + \omega_t,$$
$$c_{t+1}^* = -c_t \sin \lambda_c + c_t^* \cos \lambda_c + \omega_t^*,$$

with ω_t and ω_t^* being independent variables from $N(0, Q_{\text{cycle},t})$ distribution and frequency $\lambda_c = 2\pi/s$.

For non-Gaussian models the observation equation of (3) is replaced by $p(y_t|\theta_t)$, where $\theta_t = \mu_t + \gamma_t + c_t$. An additional Gaussian noise term ϵ_t can also be included in θ_t using the `SSMcustom` function (this is illustrated in Section 6.5). The general matrix formulation of structural time series can be found, for example, in Durbin and Koopman (2012, Chapter 3).

Three auxiliary functions, `SSMtrend`, `SSMcycle`, and `SSMseasonal`, for building structural time series are provided in **KFAS**. The argument `degree` of `SSMtrend` defines the degree of the polynomial component, where `1` corresponds to a local level model and `2` to a local linear trend model. Higher order polynomials can also be defined with larger values. Another important argument for `SSMtrend` is `Q`, which defines the covariance structure of the trend component. This is typically a list of $p \times p$ matrices (with p being the number of series for which the component is defined), where the first matrix corresponds to the level component (μ in (3)), the second to the slope component ν and so forth.

The function `SSMcycle` differs from `SSMtrend` only by one argument. `SSMcycle` does not have argument `degree`, but instead it has argument `period` which defines the length of the cycle c_t. The same argument is also used in the function `SSMseasonal`, which contains also another important argument `sea.type`, which can be used to define whether the user wants a `dummy` or a `trigonometric` seasonal.

The example models of Sections 2.2 and 3.2 are special cases of the local linear trend model, where the variance of ζ_t is zero, and there are no seasonal or cycle components. Thus the Gaussian model of Section 2.2 can be built with **KFAS** more easily by the following code:

```
R> model_structural <- SSModel(deaths / population ~
+      SSMtrend(degree = 2, Q = list(matrix(NA), matrix(0))), H = matrix(NA))
R> fit_structural <- fitSSM(model_structural, inits = c(0, 0),
+    method = "BFGS")
R> fit_structural$model["Q"]


, , 1

          [,1] [,2]
[1,] 4.256967    0
[2,] 0.000000    0
```

Here the state equation is defined using the `SSMtrend` auxiliary function without the need for an explicit definition of the corresponding system matrices. The intercept term is automatically omitted from the right side of the formula when the `SSMtrend` component is used, in order to keep the model identifiable. Here the unknown variance parameters are set to `NA`, so the default behaviour of the `fitSSM` function can be used for the parameter estimation.

6.2. ARIMA models

Another typical time series modelling framework are ARIMA models, which are also possible to define as a state space model. The auxiliary function `SSMarima` defines the ARIMA model using vectors `ar` and `ma`, which define the autoregressive and moving average coefficients, respectively. The function assumes that all series defined by the `index` have the same coefficients. The argument `d` defines the degree of differencing, and a logical argument `stationary` defines whether stationarity (after differencing) is assumed (if not, diffuse initial states are used instead of a stationary distribution). A univariate ARIMA(p,d,q) model can be written as

$$y_t^* = \phi_1 y_{t-1}^* + \ldots + \phi_p y_{t-p}^* + \xi_t + \theta_1 \xi_{t-1} + \ldots + \theta_q \xi_{t-q},$$

where $y_t^* = \Delta^d y_t$ and $\xi_t \sim \mathrm{N}(0, \sigma^2)$. Let $r = \max(p, q+1)$. **KFAS** defines the state space representation of the ARIMA(p,d,q) model with stationary initial distribution as

$$Z^\top = \begin{pmatrix} 1_{d+1} \\ 0 \\ \vdots \\ 0 \end{pmatrix}, H = 0, T = \begin{pmatrix} U_d & 1_d^\top & 0 & \cdots & 0 \\ 0 & \phi_1 & 1 & & 0 \\ \vdots & & & \ddots & \\ \vdots & \phi_{r-1} & 0 & & 1 \\ 0 & \phi_r & 0 & \cdots & 0 \end{pmatrix}, R = \begin{pmatrix} 0_d \\ 1 \\ \theta_1 \\ \vdots \\ \theta_{r-1} \end{pmatrix},$$

$$\alpha_t = \begin{pmatrix} y_{t-1} \\ \vdots \\ \Delta^{d-1} y_{t-1} \\ y_t^* \\ \phi_2 y_{t-1}^* + \ldots + \phi_r y_{t-r+1}^* + \theta_1 \eta_t + \ldots + \theta_{r-1} \eta_{t-r+2} \\ \vdots \\ \phi_r y_{t-1}^* + \theta_{r-1} \eta_t \end{pmatrix}, Q = \sigma^2,$$

$$a_1 = \begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix}, P_{*,1} = \begin{pmatrix} 0 & 0 \\ 0 & S_r \end{pmatrix}, P_{\infty,1} = \begin{pmatrix} I_d & 0 \\ 0 & 0 \end{pmatrix}, \eta_t = \xi_{t+1}$$

where $\phi_{p+1} = \ldots = \phi_r = \theta_{q+1} = \ldots = \theta_{r-1} = 0$, 1_{d+1} is a $1 \times (d+1)$ vector of ones, U_d is a $d \times d$ upper triangular matrix of ones and S_r is the covariance matrix of stationary elements of α_1. The elements of the initial state vector α_1 which correspond to the differenced values $y_0, \ldots, \Delta^{d-1} y_0$ are treated as diffuse. The covariance matrix S_r can be computed by solving the linear equation $(I - T \otimes T)\mathrm{vec}(S_r) = \mathrm{vec}(RR^\top)$ (Durbin and Koopman 2012, p.138).

Note that the `arima` function from **stats** also uses the same state space approach for ARIMA modelling, although the handling of the intercept and possible covariates is done in a slightly different manner (see documentation of `arima` for details).

As an example, we again model the alcohol-related deaths but now use the ARIMA(0,1,1) model with drift:

```
R> drift <- 1:length(deaths)
R> model_arima <- SSModel(deaths / population ~ drift +
+      SSMarima(ma = 0, d = 1, Q = 1))
```

```
R>
R> update_model <- function(pars, model) {
+    tmp <- SSMarima(ma = pars[1], d = 1, Q = pars[2])
+    model["R", states = "arima"] <- tmp$R
+    model["Q", states = "arima"] <- tmp$Q
+    model["P1", states = "arima"] <- tmp$P1
+    model
+  }
R>
R> fit_arima <- fitSSM(model_arima, inits = c(0, 1), updatefn = update_model,
+    method = "L-BFGS-B", lower = c(-1, 0), upper = c(1, 100))
R> fit_arima$optim.out$par

[1] -0.4994891 16.9937888
```

In this case we need to supply the model updating function for `fitSSM` which updates our model definition based on the current values of the parameters we are estimating. Instead of manually altering the corresponding elements of the model, `update_model` uses `SSMarima` function for computation of relevant system matrices R, Q and P_1. The estimated values for θ_1 and σ are -0.5 and 17.

Comparing the results of our previous structural time series model and the estimated ARIMA model, we see that the estimated drift term and the log-likelihood are identical:

```
R> (out_arima <- KFS(fit_arima$model))

Smoothed values of states and standard errors at time n = 39:
        Estimate  Std. Error
drift    0.8409    0.3446
arima1  20.3008   13.1100
arima2   1.3545    1.3898
arima3   0.3031    0.6453

R> (out_structural <- KFS(fit_structural$model))

Smoothed values of states and standard errors at time n = 39:
       Estimate  Std. Error
level  54.7532    2.1705
slope   0.8409    0.3446

R> out_arima$logLik

[1] -108.9734

R> out_structural$logLik

[1] -108.9734
```

This is not suprising given the well-known connections between structural time series and ARIMA models Harvey (1989).

6.3. Linear and generalized linear models

An ordinary linear regression model

$$y_t = x_t^\top \beta + \epsilon_t, \quad t = 1, \ldots, n,$$

where $\epsilon_t \sim N(0, \sigma^2)$, can be written as a Gaussian state space model by defining $Z_t = x_t^\top$, $H_t = \sigma^2$, $R_t = Q_t = 0$ and $\alpha_t = \beta$. Assuming that the prior distribution of β is defined as diffuse, the diffuse likelihood of this state space model corresponds to a restricted maximum likelihood (REML). Then the estimate for σ^2 obtained from `fitSSM` would be the familiar unbiased REML estimate of residual variance. It is important to notice that for this simple model numerical optimization is not needed, since we can estimate σ^2 by running the Kalman filter with $H_t = 1$, which gives us

$$\hat{\sigma}^2 = \frac{1}{\sum I(F_{\infty,t} = 0)} \sum_{t=1}^{n} I(F_{\infty,t} = 0) v_t^2 / F_t,$$

which equals to the REML estimate of σ^2. The initial Kalman filter already provides correct estimates of β as a_{n+1}, and running the Kalman filter again with $H_t = \sigma^2$ also gives the covariance matrix of $\hat{\beta}$ as P_{n+1}.

The extension from a linear model to a generalized linear model is straightforward as the basic theory behind the exponential family state space modelling can be formulated from the theory of generalized linear models (GLM) and can be thought of as extension to GLMs with additional dynamic structure. The iterative process of finding the approximating Gaussian model is equivalent with the famous iterative reweighted least squares (IRLS) algorithm (McCullagh and Nelder 1989, p. 40). If the model is ordinary GLM, the final estimates of regression coefficients β and their standard errors coincide with maximum likelihood estimates obtained from ordinary GLM fitting. By adjusting the prior distribution for β we can use **KFAS** also for the Bayesian analysis of Poisson and binomial regression (as those distributions do not depend on any additional parameters such as residual variance) with Gaussian prior.

A simple (generalized) linear model can be defined using `SSModel` without any auxiliary functions by defining the regression formula in the main part of the `formula`. For example, the following code defines a Poisson GLM which is identical to the one found on the help page of `glm`:

```
R> counts <- c(18, 17, 15, 20, 10, 20, 25, 13, 12)
R> outcome <- gl(3, 1, 9)
R> treatment <- gl(3, 3)
R> model_glm1 <- SSModel(counts ~ outcome + treatment,
+    distribution = "poisson")
```

The previous model could also be defined using the auxiliary function `SSMregression`:

```
R> model_glm2 <- SSModel(counts ~ SSMregression(~ outcome + treatment),
+    distribution = "poisson")
```

If our observations are multivariate, distinct regression components are defined for each of the series. For example, if counts `counts` above were a bivariate series, then both series would have ther own regression coefficients but the same covariate values. By using `SSMregression` explicitly, one could also define `type = "common"`, which would construct common regression coefficients for all series.

With `SSMregression` one can also define more complex regression models. The first argument of `SSMregression`, `rformula` can be used to provide a single formula or a list of formulas, where each component of the list contains the appropriate formula to be used for the corresponding series (ith formula in the list is used for the ith series defined by the argument `index`). When `rformula` is a list, the `data` argument of `SSMregression` can be a single data frame (or environment), or a list of such data objects. If `data` is a list, ith element of that list is used for ith formula, and if `data` is a single data frame or environment, the same data is used for all formulas.

The state space approach makes it possible to extend classical GLMs in many ways. The extension to multivariate GLMs is straightforward, allowing, for example, the modelling of multiple groups of data where some of the model parameters are assumed to be identical between the groups, or where the number of explanatory variables differs between groups. For Gaussian models, a correlation of error terms ϵ between groups can also be incorporated. The use of dynamic GLM, where the regression coefficients follow a random walk process, can be defined by using argument `Q` in `SSMregression`. By manually altering the corresponding elements in the T matrix, one can also define autoregressive behaviour for the coefficients. An additional parameter u_t is defined separately for each observation, making it possible to define models where, for example, the dispersion parameter of a negative binomial model varies in time. One can also compute prediction intervals and other interesting measures efficiently via the importance sampling approach discussed in Section 3.

6.4. Generalized linear mixed models

Just like in GLM setting, it is also possible to write the generalized linear mixed model (GLMM) as a state space model. The difference between fixed and random effects lies in the initial state distribution; fixed effects are initialized via diffuse prior whereas random effects have proper variance defined by elements of P_1. Both types of states are automatically estimated by the Kalman filter, given the covariance structure of the random effects (and the residual variance or other parameters related to the distribution of the observation equation).

In practice, the mixed model formulation becomes quite cumbersome especially in hierarchical settings, but with large longitudinal settings it might still be useful to write a mixed model as a state space model, as it is then straightforward to add, for example, stochastic cycles or trends to the model. As an example I define a linear mixed model for the sleep deprivation study data from **lme4** (Bates, Mächler, Bolker, and Walker 2015a; Bates, Maechler, Bolker, and Walker 2015b) package as on the help page of the data. The data frame `sleepstudy` consists of three variables, the response variable `Reaction` (average reaction time), `Days` (number of days of sleep deprivation) and grouping variable `Subject`. First the response variable is restructured to a `matrix` (or `ts`) object:

```
R> library("lme4", quietly = TRUE)
R> y_split <- split(sleepstudy["Reaction"], sleepstudy["Subject"])
R> p <- length(y_split)
```

```
R> y <- matrix(unlist(y_split), ncol = p,
+    dimnames = list(NULL, paste("Subject", names(y_split))))
```

The data frame with explanatory variables is also split to a list where each list component corresponds to one group.

```
R> dataf <- split(sleepstudy, sleepstudy["Subject"])
```

The only explanatory variable `Days` in the data is identical to each Subject so the previous split of the data frame is not necessary, but illustrates the workflow for more complex data.

We can now build the state space model by defining the common fixed part for each group (`SSMregression` function with argument `type = "common"`). Using the same function we can define the distinct random effect parts for each group, and the covariance structure of the random effects using the argument `P1` (the diffuse part `P1inf` is automatically set to zero for those states where the corresponding element in `P1` is nonzero). The function `.bdiag` from the **Matrix** package (Bates and Maechler 2015) is used for building a block diagonal covariance matrix for the random effects.

```
R> P1 <- as.matrix(.bdiag(replicate(p, matrix(NA, 2, 2), simplify = FALSE)))
R> model_lmm <- SSModel(y ~ -1 +
+      SSMregression(rep(list(~ Days), p), type = "common", data = dataf,
+        remove.intercept = FALSE) +
+      SSMregression(rep(list(~ Days), p), data = dataf,
+        remove.intercept = FALSE, P1 = P1),
+    H = diag(NA, p))
```

Note that in `SSMregression`, we use a lists of formulas (the first unnamed argument `rformula`) and data frames (argument `data`), where each component corresponds to one column of `y`. Thus we could use different formulas for different groups in more complex models. In this simple example the same model could be built with call

```
R> model_lmm2 <- SSModel(y ~ - 1 +
+      SSMregression(~ Days, type = "common", remove.intercept = FALSE) +
+      SSMregression(~ Days, remove.intercept = FALSE, P1 = P1),
+    H = diag(NA, p), data = data.frame(Days = 0:9))
```

Again we need to define the model updating function for `fitSSM`:

```
R> update_lmm <- function(pars, model) {
+    P1 <- diag(exp(pars[1:2]))
+    P1[1, 2] <- pars[3]
+    P1 <- crossprod(P1)
+    model["P1", states = 3:38] <-
+      as.matrix(.bdiag(replicate(p, P1, simplify = FALSE)))
+    model["H"] <- diag(exp(pars[4]), p)
+    model
+  }
R>
R> fit_lmm <- fitSSM(model_lmm, c(1, 1, 1, 5), update_lmm, method = "BFGS")
```

The estimated likelihood, variance/covariance parameters, and the estimates of fixed and random effects are practically identical to the ones obtained by **lmer** function of **lme4** package. The only major difference is in the estimation of conditional covariance matrices of the random effects, which is due to the fact that in **lme4** these matrices are computed conditionally on all the other model parameters, including the fixed effects (Bates *et al.* 2015a, p.28). In **KFAS** the conditioning is only on the the numerically estimated variance/covariance parameters, and thus the resulting standard errors of random effects take account of the uncertainty of the estimation of fixed effects also.

In this example different groups were thought of as separate response vectors. In cases where the sample sizes in different groups are not equal, the same approach can be used after appropriately filling the data matrix with missing values. The corresponding `NA` values in covariates do not cause problems as they are not referenced in the Kalman filter.

It is also possible to define the mixed model using univariate response and time-varying system matrices T_t and Q_t. This reduces the state space and thus makes the model computationally more efficient, but adding other stochastic components to the model can be more problematic. For building such a model we need to use either the `customSSM` function for defining the corresponding time-varying system matrices, or we can use `SSMregression` as a starting point and alter the model components manually. This univariate approach is illustated on the main help page of **KFAS**.

6.5. Arbitrary state space models

By combining the auxialiary functions presented in the previous sections and possibly manually adjusting the resulting system matrices, a large amount of models can be constructed with relative ease. For cases where this is not sufficient or otherwise preferable, the auxiliary function `SSMcustom` can be used for the construction of arbitrary components by direct definition of the system matrices. As an example, we modify the Poisson model of Section 3 by adding an additional white noise term which tries to capture possible overdispersion of the data. Our model for the Poisson intensity is now $u_t \exp(\mu_t + \epsilon_t)$ with

$$\mu_{t+1} = \mu_t + \nu + \eta_t,$$

where $\eta_t \sim N(0, \sigma_\eta^2)$ as before, and $\epsilon_t \sim N(0, \sigma_\epsilon^2)$. This model can be written in a state space form by defining

```
R> model_poisson <- SSModel(deaths ~ SSMtrend(2, Q = list(NA, 0)) +
+      SSMcustom(Z = 1, T = 0, Q = NA, P1 = NA),
+    distribution = "poisson",  u = population)
```

As the model contains unknown parameters in `P1`, we need to provide a specific model updating function for `fitSSM`:

```
R> update_poisson <- function(pars, model) {
+    model["Q", etas = "level"] <- exp(pars[1])
+    model["Q", etas = "custom"] <- exp(pars[2])
+    model["P1", states = "custom"] <- exp(pars[2])
+    model
+  }
```

```
R> fit_poisson <- fitSSM(model_poisson, c(-3, -3),
+    update_poisson, method = "BFGS")
R> fit_poisson$model["Q", etas = "level"]

[1] 0.00316852

R> fit_poisson$model["Q", etas = "custom"]

[1] 0.002506342
```

From Figure 3 we see that the Gaussian structural time series model and the Poisson structural time series model with additional white noise produce nearly indistinguishable estimates of the smoothed trend μ_t. This is due to the relatively high intensity of the Poisson process.

Figure 3: Alcohol-related deaths in Finland (black line) with smoothed estimates from Gaussian model (blue) and Poisson model with additional noise (red).

7. Illustration

I now illustrate the use of **KFAS** with a more complete example case than the previous examples. Again the data consists of alcohol-related deaths in Finland, but now four age groups, 30–39, 40–49, 50–59 and 60–69, are modelled together as a multivariate Poisson model.

The death counts and yearly population sizes in corresponding age groups are available for the years 1969–2012, but as an illustration, we only use the data until 2007, and make predictions for the years 2008–2013. Figure 4 shows the number of deaths per 100,000 persons for all age groups.

```
R> data("alcohol")
R> colnames(alcohol)
```

```
[1] "death at age 30-39"      "death at age 40-49"
[3] "death at age 50-59"      "death at age 60-69"
[5] "population by age 30-39" "population by age 40-49"
[7] "population by age 50-59" "population by age 60-69"
```

```
R> ts.plot(window(alcohol[, 1:4] / alcohol[, 5:8], end = 2007), col = 1:4,
+    ylab = "Alcohol-related deaths in Finland per 100,000 persons",
+    xlab = "Year")
R> legend("topleft",col = 1:4, lty = 1, legend = colnames(alcohol)[1:4])
```

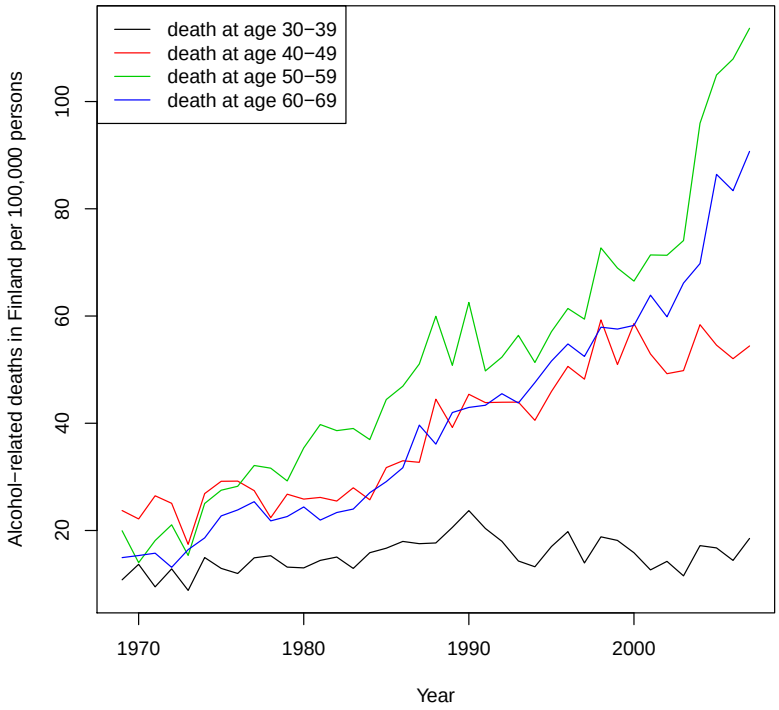

Figure 4: Alcohol-related deaths per 100,000 persons in Finland in 1969–2007 for four age groups.

Here I choose a multivariate extension of the Poisson model used in Section 6.5:

$$\begin{aligned} p(y_t|\theta_t) &= \text{Poisson}(u_t e^{\theta_t}), \quad u_t = \text{population}_t, \\ \theta_t &= \mu_t + \epsilon_t, \quad \epsilon_t \sim N(0, Q_{\text{noise}}), \\ \mu_{t+1} &= \mu_t + \nu_t + \xi_t, \quad \xi_t \sim N(0, Q_{\text{level}}), \\ \nu_{t+1} &= \nu_t. \end{aligned} \tag{4}$$

Here μ_t is the random walk with drift component, ν_t is a constant slope and ϵ_t is an additional white noise component which captures the extra variation of the series. I make no restrictions for the covariance structures of the level or the noise component.

The model (4) can be constructed with **KFAS** as follows.

```
R> alcoholPred <- window(alcohol, start = 1969, end = 2007)
R> model <- SSModel(alcoholPred[, 1:4] ~
+      SSMtrend(2, Q = list(matrix(NA, 4, 4), matrix(0, 4, 4))) +
+      SSMcustom(Z = diag(1, 4), T = diag(0, 4), Q = matrix(NA, 4, 4),
+        P1 = matrix(NA, 4, 4)), distribution = "poisson",
+    u = alcoholPred[, 5:8])
```

The updating function for `fitSSM` is

```
R> updatefn <- function(pars, model, ...){
+    Q <- diag(exp(pars[1:4]))
+    Q[upper.tri(Q)] <- pars[5:10]
+    model["Q", etas = "level"] <- crossprod(Q)
+    Q <- diag(exp(pars[11:14]))
+    Q[upper.tri(Q)] <- pars[15:20]
+    model["Q", etas = "custom"] <- model["P1", states = "custom"] <-
+      crossprod(Q)
+    model
+  }
```

We can estimate the model parameters first without simulation, and then using those estimates as initial values run the estimation procedure again with importance sampling. In this case, the results obtained from the importance sampling step are practically identical with the ones obtained from the initial step.

```
R> init <- chol(cov(log(alcoholPred[, 1:4] / alcoholPred[, 5:8])) / 10)
R> fitinit <- fitSSM(model, updatefn = updatefn,
+    inits = rep(c(log(diag(init)), init[upper.tri(init)]), 2),
+    method = "BFGS")
R> -fitinit$optim.out$val

[1] -704.8052

R> fit <- fitSSM(model, updatefn = updatefn, inits = fitinit$optim.out$par,
+    method = "BFGS", nsim = 250)
R> -fit$optim.out$val
```

```
[1] -704.8034
```

Using the model extraction method for the fitted models, we can check the estimated covariance and correlation matrices:

```
R> varcor <- fit$model["Q", etas = "level"]
R> varcor[upper.tri(varcor)] <- cov2cor(varcor)[upper.tri(varcor)]
R> print(varcor, digits = 2)
```

```
       [,1]    [,2]   [,3]  [,4]
[1,] 0.0074 0.66022 0.8062 0.856
[2,] 0.0028 0.00239 0.1654 0.711
[3,] 0.0040 0.00047 0.0034 0.755
[4,] 0.0033 0.00156 0.0020 0.002
```

```
R> varcor <- fit$model["Q", etas = "custom"]
R> varcor[upper.tri(varcor)] <- cov2cor(varcor)[upper.tri(varcor)]
R> print(varcor, digits = 2)
```

```
        [,1]    [,2]    [,3]    [,4]
[1,] 0.00537 0.73118 0.75627 8.0e-01
[2,] 0.00315 0.00346 0.99924 9.9e-01
[3,] 0.00295 0.00313 0.00283 1.0e+00
[4,] 0.00043 0.00043 0.00039 5.4e-05
```

Parameter estimation of a state space model is often a difficult task, as the likelihood surface contains multiple maxima, thus making the optimization problem highly dependent on the initial values. Often the unknown parameters are related to the unobserved latent states, such as the covariance matrix in this example, with little a priori knowledge. Therefore, it is challenging to guess good initial values, especially in more complex settings. Thus, multiple initial value configurations possibly with several different type of optimization routines is recommended before one can be reasonably sure that proper optimum is found. Here we use the covariance matrix of the observed series as initial values for the covariance structures.

Another issue in the case of non-Gaussian models is the fact that the likelihood computation is based on iterative procedure which is stopped using some stopping criteria (such as the relative change of log-likelihood), so the log-likelihood function actually contains some noise. This in turn can affect the gradient computations in methods like BFGS and can in theory give unreliable results. Using derivative free method like Nelder-Mead is therefore sometimes recommended. On the other hand, BFGS is usually much faster than Nelder-Mead, and thus I prefer to try BFGS first at least in preliminary analysis.

Using the function `KFS` we can compute the smoothed estimates of states:

```
R> out <- KFS(fit$model, nsim = 1000)
R> out
```

```
Smoothed values of states and standard errors at time n = 39:
                          Estimate    Std. Error
level.death at age 30-39   2.8559160   0.0784371
slope.death at age 30-39   0.0107142   0.0137135
level.death at age 40-49   4.0313117   0.0423763
slope.death at age 40-49   0.0237188   0.0076318
level.death at age 50-59   4.7578026   0.0398295
slope.death at age 50-59   0.0503715   0.0095850
level.death at age 60-69   4.4938371   0.0332897
slope.death at age 60-69   0.0482386   0.0072090
custom1                   -0.0004021   0.0603946
custom2                   -0.0195488   0.0408846
custom3                   -0.0169493   0.0370236
custom4                   -0.0021345   0.0051427
```

From the output of `KFS` we see that the slope term is not significant in the first age group. For time-varying states we can easily plot the estimated level and noise components, which shows clear trends in three age groups and highly correlated additional variation in all groups:

```
R> plot(coef(out, states = c("level", "custom")), main = "Smoothed states",
+    yax.flip = TRUE)
```

Note the large drop in the noise component in Figure 5, which relates to a possible outlier in 1973 of the mortality series. As an illustration of model diagnostics, we compute recursive residuals for our model and check whether there is autocorrelation left in the residuals (Figure 6).

```
R> res <- rstandard(KFS(fit$model, filtering = "mean", smoothing = "none",
+    nsim = 1000))
R> acf(res, na.action = na.pass)
```

We see occasional lagged cross-correlation between the residuals, but overall we can be relatively satisfied with our model.

We can now predict the intensity e^{θ_t} of alcohol-related deaths per 100,000 persons for each age group for years 2008–2013 using our estimated model. As our model is time varying (`u` varies), we need to provide the model for the future observations via `newdata` argument. In this case we can use the `SSMcustom` function and provide all the necessary system matrices at once, together with constant `u = 1` (our signal θ is already scaled properly as the original u_t was the population per 100,000 persons).

```
R> pred <- predict(fit$model,
+    newdata = SSModel(ts(matrix(NA, 6, 4), start = 2008) ~ -1 +
+        SSMcustom(Z = fit$model$Z, T = fit$model$T, R = fit$model$R,
+          Q = fit$model$Q), u = 1, distribution = "poisson"),
+    interval = "confidence", nsim = 10000)
```

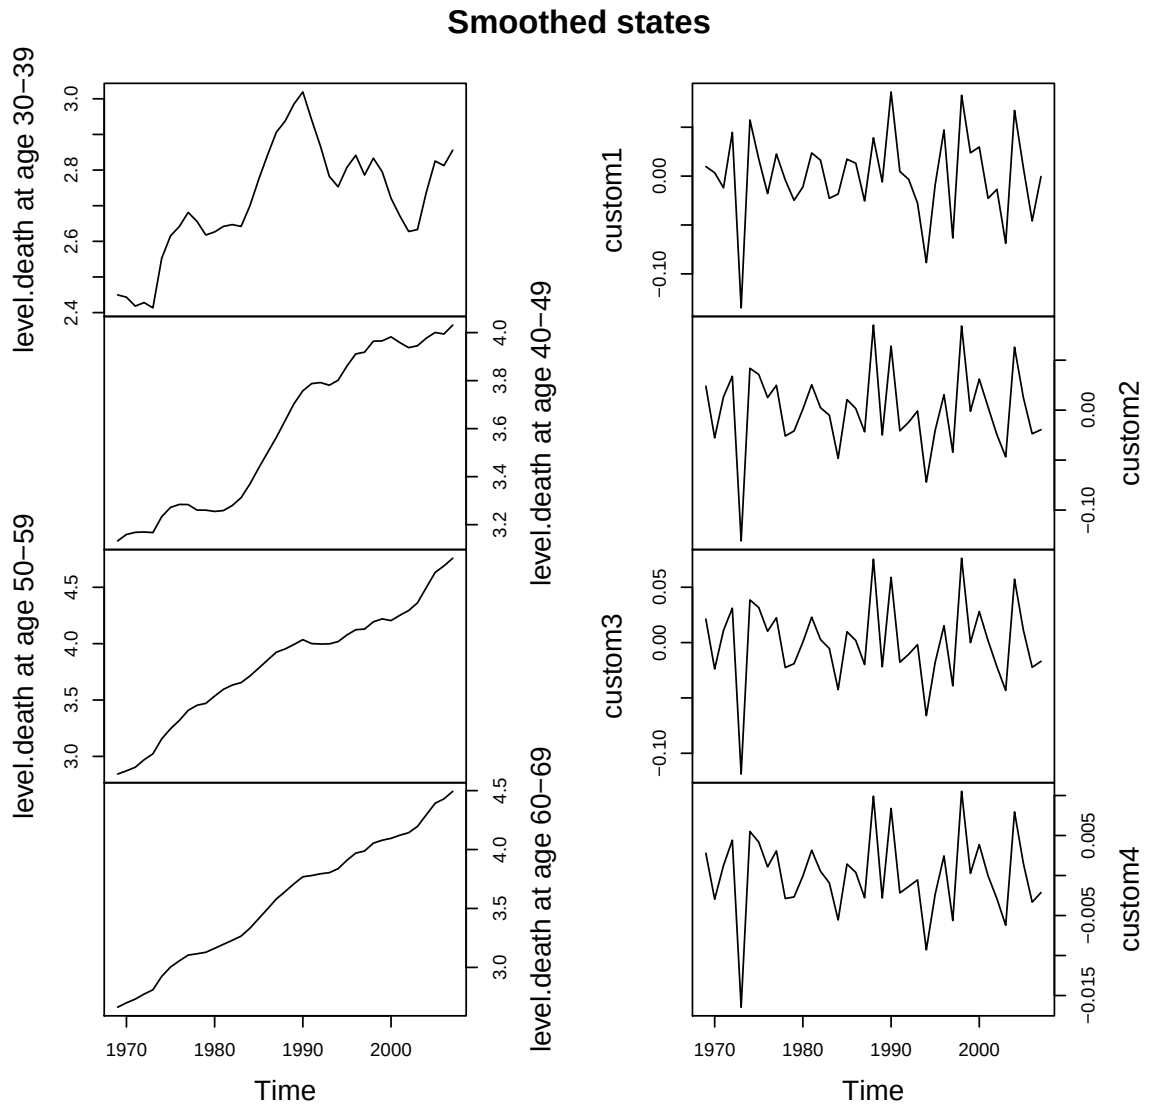

Figure 5: Smoothed level and white noise components.

```
R> trend <- exp(signal(out, "trend")$signal)
R> par(mfrow = c(2, 2), mar = c(2, 2, 2, 2) + 0.1, oma = c(2, 2, 0, 0))
R> for (i in 1:4) {
+    ts.plot(alcohol[, i]/alcohol[, 4 + i], trend[, i], pred[[i]],
+      col = c(1, 2, rep(3, 3)), xlab = NULL, ylab = NULL,
+      main = colnames(alcohol)[i])
+  }
R> mtext("Number of alcohol related deaths per 100,000 persons in Finland",
+    side = 2, outer = TRUE)
R> mtext("Year", side = 1, outer = TRUE)
```

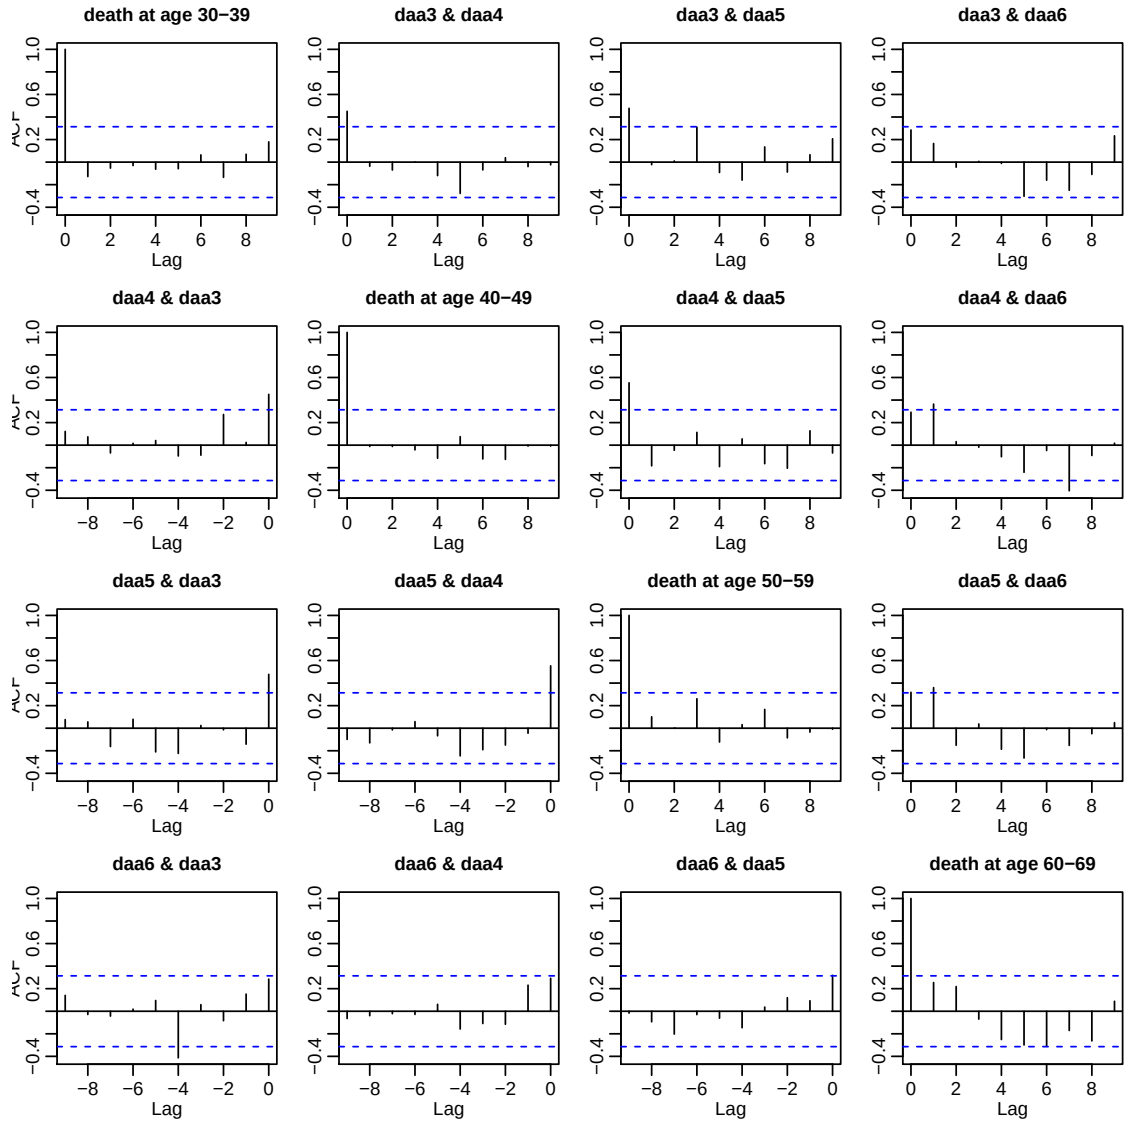

Figure 6: Autocorrelations and cross-correlations of recursive residuals.

Figure 7 shows the observed deaths, smoothed trends for 1969–2007, and intensity predictions for 2008–2013 together with 95% prediction intervals for intensity. When we compare our predictions with true observations, we see that in reality the number of deaths slightly increased in the oldest age group (ages 60–69), whereas at another age they decreased substantially during the forecasting period. This is partly explained by the fact that during this period the total alcohol consumption decreased almost monotonically, which in turn might have been caused by the increase in taxation of alcohol in 2008, 2009 and 2012.

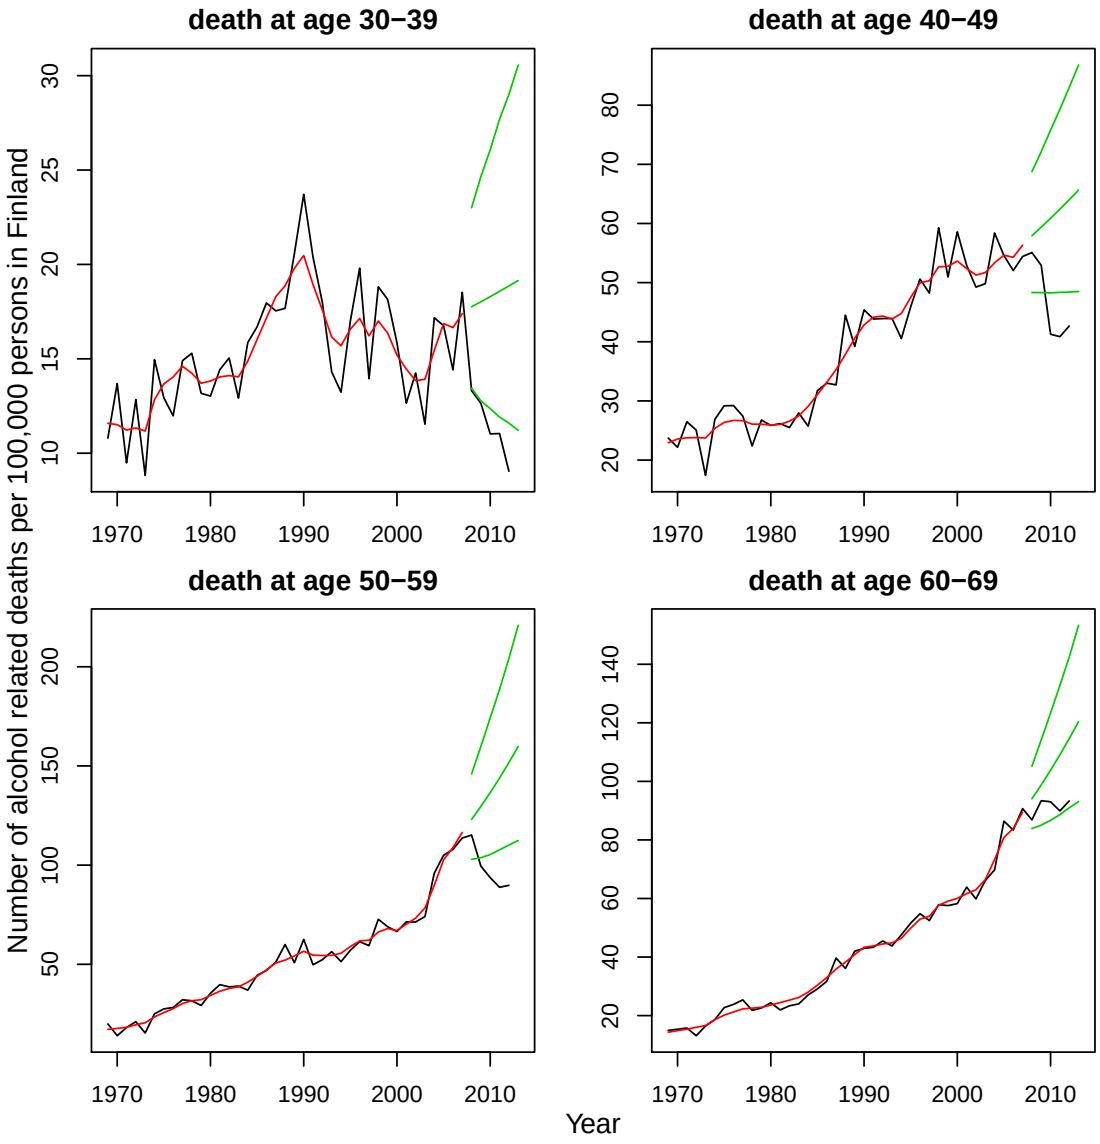

Figure 7: Observed number of alcohol related deaths per 100,000 persons in Finland (black), fitted values (red) and intensity predictions for years the 2008–2013 together with 95% prediction intervals (green).

8. Other packages for non-Gaussian time series modelling

There are also other packages in CRAN which can be used for modelling non-Gaussian time series data. Package **pomp** (King, Ionides, Bretó, Ellner, Ferrari, Kendall, Lavine, Nguyen, Reuman, Wearing, and Wood 2015a; King, Nguyen, and Ionides 2015b) offers functions for the inference of state space models with non-Gaussian and non-linear observation and state equations via particle filtering methods. The particle filtering approach makes **pomp** applicable to an even more broader class of models than **KFAS**, but the learning curve for using **pomp** is relatively high, as the user must write his or her own functions (preferably in C) for measurement and state process simulation as well as likelihood evaluation. Another package suitable for state space modelling is **INLA** (Rue, Martino, Lindgren, Simpson, Riebler, and Krainski 2015; Lindgren and Rue 2015) (not available on CRAN), which can be used for Bayesian analysis via integrated nested Laplace approximation technique. Although it is often used in spatial modelling via Gaussian random fields, it can also be used for certain temporal state space models where the state transitions are Gaussian.

KFAS is based on a parameter-driven approach where the latent states α_t evolve in time as stochastic processes with the noise term η_t which does not depend on the past observations or covariates. This approach offers a flexible and conceptually simple way of introducing multiple types of latent structures into the model. In contrast, in the observation-driven approach the state equation is defined using the past observations and possibly other covariates. This makes the states perfectly predictable (one-step-ahead) given the past information, allowing closed-form evaluation of the likelihood, and thus leading to computational gains compared to simulation-based estimation methods used in the parameter-driven approach for non-Gaussian models. Both approaches have their merits, see, for example, (Koopman, Lucas, and Scharth 2015) for a comparison of parameter-driven and observation-driven approaches in a complex non-Gaussian non-linear setting.

acp (Vasileios 2015) is a compact package based on the observation-driven approach for count data regression via Autoregressive Conditional Poisson (ACP) processes. In the ACP models the mean of the Poisson process is assumed to depend on the previous values of the observations and the previous values of the mean (which in turn can depend on covariates). A more general framework to the observation-driven approach for time series regression is implemented in the package **glarma** (Dunsmuir and Scott 2015), which implements Generalized Linear Autoregressive Moving Average models (GLARMA) supporting Poisson, binomial and negative binomial distributions. Package **tscount** (Liboschik, Fried, Fokianos, and Probst 2015) offers similar functionality using Poisson and negative binomial distributions with a closely related theoretic framework. All of these three packages assume univariate responses.

The scope of the packages **gamlls** (Rigby and Stasinopoulos 2005) and **VGAM** (Yee 2010) is mainly on complex non-time series data, but they also have some capabilities for non-Gaussian time series modelling. Package **gamlss.util** (Stasinopoulos, Rigby, and Eilers 2015) extends **gamlls** with the function `garmaFit` for univariate time series regression via GARMA models, which are closely related to the GLARMA models of the **glarma** package. A large number of distributions are supported. The GARMA models can also be estimated with **VGAM** (Yee 2010) package, which contains the function `garma` for the estimation of GARMA models, but the documentation warns that the function is very unpolished.

Time series of counts often exhibit overdispersion or an excess amount of zeroes. Although previously mentioned packages can deal with these issues to some extent, there are also

packages on CRAN designed specifically for these type of problems. Package **ZIM** (Yang, Zamba, and Cavanaugh 2014) offers functions for both observation-driven and parameter-driven modelling of zero-inflated count series. For the parameter-driven approach a particle filtering approach is used. From the very scarce documentation of the package, it is not clear how the observation-driven approach is implemented. Package **tsintermittent** (Kourentzes and Petropoulos 2015) contain forecasting methods for intermittent time series stemming, for example, from sales of slow moving items. Covariates are not supported.

Overall, there are multiple packages on CRAN which offer different approaches to non-Gaussian time series modelling, and preferring one package over another is likely dependent on the current problem in hand. For example, in some cases, time-dependency in the data can be more thought of as a nuisance which must be taken into account in order to make reliable inferences regarding the regression coefficients of the model. However, in some cases it can be that the interest is in the underlying latent time-varying processes itself. Due to the parameter-driven approach, packages such as **KFAS** and **pomp** can be used for the flexible modelling of, for example, a stochastic trend, seasonal and cyclic components. For packages such as **glarma** and **tscount** options are more limited because of the nature of the general model specification.

Time-varying regression coefficients and random effects can be incorporated to time series models with **INLA**, **KFAS** and **pomp**. These packages can also deal with missing observations in the response variable straightforwardly, whereas other packages do not seem to handle missing values properly. Most of the packages produce informative or non-informative error messages in the case of missing observations, whereas some just omit the missing time points of the data without taking account of the unevenness of the time points during the parameter estimation.

8.1. Comparison to INLA

I will now briefly compare **KFAS** and **INLA**. As an illustration, we reanalyze the salmonella data analyzed by Margolin, Kaplan, and Zeiger (1981) which is available from **INLA**. The data consists of the number of revertant colonies of TA98 Salmonella with different doses of quinoline. We model the number of colonies as a Poisson GLMM with two explanatory variables and a random intercept term which tries to capture the overdispersion in the data. The codes for inference with **INLA** used here can be found from `http://www.r-inla.org/examples/volume-1/code-for-salm-example`. The **INLA** is not available at CRAN but can be downloaded from `http://www.math.ntnu.no/inla/R/stable`.

```
R> library("INLA", quietly = TRUE)
R> data("Salm")
R> mod.salm <- inla(y ~ log(dose + 10) + dose +
+      f(rand, model = "iid", param = c(0.001, 0.001)),
+    family = "poisson", data = Salm)
R> h.salm <- inla.hyperpar(mod.salm)
```

There are two ways to define the random intercept component in **KFAS**. The first one uses the `SSMregression` function and constructs a factor with 18 levels (one for each case) with non-diffuse initial variance σ^2. This gives 18 identically distributed time-invariant states, where each state corresponds to the random effect of one observation. Another option would be

to use the `SSMcustom` function and define just one time-varying state as `SSMcustom(Z = 1, T = 0, R = 1, Q = sigma2, a1 = 0, P1 = sigma2, P1inf = 0)`. Both approaches give identical results. However, for large data the former approach is less efficient as the number of states depends on the number of observations. Nevertheless, we use the former approach here for illustration.

```
R> Salm$rand <- as.factor(Salm$rand)
R> model <- SSModel(y ~ log(dose + 10) + dose +
+      SSMregression(~ -1 + rand, P1 = diag(NA, 18),
+        remove.intercept = FALSE),
+    data = Salm, distribution = "poisson")
R>
R> updatefn <- function(pars,model,...){
+    diag(model["P1", states = 4:21]) <- exp(pars)
+    model
+  }
R>
R> fit <- fitSSM(model, updatefn = updatefn, inits = -3, method = "BFGS",
+    nsim = 1000)

R> out <- KFS(fit$model, nsim = 10000)
R> out

Smoothed values of states and standard errors at time n = 18:
                Estimate    Std. Error
(Intercept)      2.1674035   0.3519119
log(dose + 10)   0.3123230   0.0957658
dose            -0.0009783   0.0004249
rand1           -0.0998189   0.1929948
rand2            0.0811525   0.1905155
rand3            0.2962272   0.1852740
rand4           -0.1796443   0.1826173
rand5           -0.1210575   0.1804227
rand6           -0.0365042   0.1757863
rand7           -0.2997077   0.1738915
rand8           -0.0405697   0.1641720
rand9            0.1189599   0.1591311
rand10          -0.1649685   0.1664616
rand11           0.1202154   0.1569226
rand12           0.4328667   0.1490611
rand13          -0.1235595   0.1659587
rand14          -0.0274907   0.1650945
rand15           0.0274298   0.1624904
rand16          -0.2166174   0.1950416
rand17          -0.0434169   0.1915480
rand18           0.2725027   0.1914863
```

```
R> h.salm$summary.fixed[, 1:2]

                       mean           sd
(Intercept)     2.1671787697 0.3627732056
log(dose + 10)  0.3127852007 0.0987778011
dose           -0.0009794203 0.0004366782

R> h.salm$summary.random$rand[, 2:3]

          mean        sd
1  -0.09491396 0.1979799
2   0.08421357 0.1929301
3   0.29577465 0.1946343
4  -0.17165618 0.1860174
5  -0.11379450 0.1819493
6  -0.03076457 0.1776603
7  -0.29145696 0.1842334
8  -0.03433952 0.1653510
9   0.12229514 0.1607173
10 -0.15831766 0.1695020
11  0.12169606 0.1603516
12  0.42740226 0.1589524
13 -0.11885975 0.1685411
14 -0.02364225 0.1653818
15  0.03042788 0.1641417
16 -0.21238483 0.2025258
17 -0.04120592 0.1964929
18  0.26895855 0.1967189

R> 1 / h.salm$summary.hyper[1]

                          mean
Precision for rand 0.04825638

R> fit$model["P1", states = 4]

[1] 0.06554971
```

Although **INLA** uses a Bayesian approach, which takes account of the parameter estimation uncertainty, the results from **INLA** and **KFAS** are practically the same, even with such small data. The Kalman filtering with diffuse initialization still takes account of the uncertainty of the estimation of regression coefficients, so the differences here are related to the different prior definitions and the estimation of the hyperparameter σ^2, which is estimated as precision $1/\sigma^2$ by **INLA**. The estimate of σ^2 by **KFAS** is 0.066 whereas **INLA** gives $\sigma^2 = 0.048$. Note that changing the estimated σ^2 from **INLA** into the model estimated by **KFAS** produces a slightly lower log-likelihood value (-73.50 versus -73.66).

As **INLA** and **KFAS** are based on a different (although related) theoretical framework, the extensive study of their performances in terms of the computational efficiency and accuracy of results is somewhat pointless. Nevertheless, some remarks can be made. I feel that the biggest advantage of **INLA** is the Bayesian framework which allows us to take account of the parameter uncertainty in predictions and other inference. On the other hand, the computational burden related to the numerical integration over the hyperparameters can become infeasible as the number of hyperparameters increases. It is not uncommon to have a time series model with tens (or even hundreds) of parameters (such as multivariate structural time series or dynamic factor models). Of course, these same models can cause problems also to the maximum likelihood estimation, as noted in the Section 7. Also the Bayesian approach eliminates the need for defining good initial values for the maximum likelihood estimation but the problem transforms into defining good priors for the same hyperparameters, which again is a non-trivial task in practice.

9. Discussion

State space models offer tools for solving a large class of statistical problems. Here I introduced an R package **KFAS** for linear state space modelling where the observations are from an exponential family. With such a general framework, different aspects of the modelling need to be taken into account. Therefore the focus of the package has been to provide reliable and relatively fast tools for multiple inference problems, such as maximum likelihood estimation, filtering, smoothing and simulation. Compared with the early versions of **KFAS**, constructing a state space model with simple components is now possible without an explicit definition of the system matrices by using the auxiliary functions and symbolic descriptions with the help of formula objects, which should greatly ease the use of the package.

Currently all the time consuming parts of **KFAS** are written in Fortran, which makes it relatively fast, given the general nature of the problems **KFAS** can handle. Still, converting the package to C++ and `S4` classes with the help of **Rcpp** (Eddelbuettel and François 2011; Eddelbuettel 2013) could result in potential improvements in terms of memory management, scalability and maintenance.

Acknowledgments

The author wishes to thank Jukka Nyblom, Patricia Menendez, Spencer Graves, as well as the editor and two anonymous reviewers for the valuable comments and suggestions regarding the paper and the package. Comments, suggestions and bug reports from various users of **KFAS** over the years are also highly appreciated. The author is also grateful for the financial support from the Emil Aaltonen Foundation.

A. Appendix: Filtering and smoothing recursions

The following formulas summarize the Kalman filtering and smoothing formulas for diffuse and sequential case andare based on Durbin and Koopman (2012) and related articles. The original formulas are somewhat scattered between the references with slightly different notations. Therefore I have collected the equations used in **KFAS** to this Appendix.

A.1. Filtering

Denote

$$a_{t+1} = \mathsf{E}(\alpha_{t+1}|y_t, \ldots, y_1) \quad \text{and}$$
$$P_{t+1} = \mathsf{VAR}(\alpha_{t+1}|y_t, \ldots, y_1).$$

The Kalman filter recursions for the general Gaussian model of form (1) are

$$\begin{aligned} v_t &= y_t - Z_t a_t \\ F_t &= Z_t P_t Z_t^\top + H_t \\ K_t &= P_t Z_t^\top \\ a_{t+1} &= T_t(a_t + K_t F_t^{-1} v_t) \\ P_{t+1} &= T_t(P_t - K_t F_t^{-1} K_t^\top) T_t^\top + R_t Q_t R_t, \end{aligned}$$

For the univariate approach, the filtering equations are

$$\begin{aligned} v_{t,i} &= y_{t,i} - Z_{t,i} a_{t,i} \\ F_{t,i} &= Z_{t,i} P_{t,i} Z_{t,i}^\top + \sigma^2_{t,i} \\ K_{t,i} &= P_{t,i} Z_{t,i}^\top \\ a_{t,i+1} &= a_{t,i} + K_{t,i} F_{t,i}^{-1} v_{t,i} \\ P_{t,i+1} &= P_{t,i} - K_{t,i} K_{t,i}^\top F_{t,i}^{-1} \\ a_{t+1,1} &= T_t a_{t,p_t+1} \\ P_{t+1,1} &= T_t P_{t,p_t+1} T_t^\top + R_t Q_t R_t, \end{aligned}$$

for $t = 1, \ldots, n$ and $i = 1, \ldots, p_t$, where $v_{t,i}$ and $F_{t,i}$ are scalars, $K_{t,i}$ is a column vector and $\sigma^2_{t,i}$ is the ith diagonal element of H_t. It is possible that $F_{t,i} = 0$, which case $a_{t,i+1} = a_{t,i}$, $P_{t,i+1} = P_{t,i}$, and $v_{t,i}$ is computed as usual.

The diffuse filtering equations for univariate approach are

$$\begin{aligned} v_{t,i} &= y_{t,i} - Z_{t,i} a_{t,i} \\ F_{*,t,i} &= Z_{t,i} P_{*,t,i} Z_{t,i}^\top + \sigma^2_{t,i} \\ F_{\infty,t,i} &= Z_{t,i} P_{\infty,t,i} Z_{t,i}^\top \\ K_{*,t,i} &= P_{*,t,i} Z_{t,i}^\top \\ K_{\infty,t,i} &= P_{\infty,t,i} Z_{t,i}^\top, \end{aligned}$$

and

$$\begin{aligned} a_{t,i+1} &= a_{t,i} + K_{\infty,t,i} v_{t,i} F_{\infty,t,i}^{-1} \\ P_{*,t,i+1} &= P_{*,t,i} + K_{\infty,t,i} K_{\infty,t,i}^\top F_{*,t,i} F_{\infty,t,i}^{-2} - (K_{*,t,i} K_{\infty,t,i}^\top + K_{*,t,i} K_{\infty,t,i}^\top) F_{\infty,t,i}^{-1} \\ P_{\infty,t,i+1} &= P_{\infty,t,i} - K_{\infty,t,i} K_{\infty,t,i}^\top F_{\infty,t,i}^{-1} \end{aligned}$$

if $F_{\infty,t,i} > 0$, and

$$\begin{aligned} a_{t,i+1} &= a_{t,i} + K_{*,t,i} v_{t,i} F_{*,t,i}^{-1} \\ P_{*,t,i+1} &= P_{*,t,i} - K_{*,t,i} K_{*,t,i}^\top F_{*,t,i}^{-1} \\ P_{\infty,t,i+1} &= P_{\infty,t,i}, \end{aligned}$$

if $F_{\infty,t,i} = 0$. The transition equations from t to $t+1$ are

$$\begin{aligned} a_{t+1,1} &= T_t a_{t,p_t+1} \\ P_{*,t+1,1} &= T_t P_{*,t,p_t+1} T_t^\top + R_t Q_t R_t \\ P_{\infty,t+1,1} &= T_t P_{\infty,t,p_t+1} T_t^\top. \end{aligned}$$

A.2. Smoothing

Denote

$$\begin{aligned} \hat{\alpha}_t &= \mathsf{E}(\alpha_t | y_n, \ldots, y_1) \quad \text{and} \\ V_t &= \mathsf{VAR}(\alpha_t | y_n, \ldots, y_1). \end{aligned}$$

The smoothing algorithms of **KFAS** are based on the following recursions:

$$\begin{aligned} r_{t,i-1} &= Z_{t,i}^\top v_{t,i} F_{t,i}^{-1} + L_{t,i}^\top r_{t,i}, \\ r_{t-1,p_t} &= T_{t-1}^\top r_{t,0}, \\ N_{t,i-1} &= Z_{t,i}^\top Z_{t,i} F_{t,i}^{-1} + L_{t,i}^\top N_{t,i} L_{t,i}, \\ N_{t-1,p_t} &= T_{t-1}^\top N_{t,0} T_{t-1}, \\ L_{t,i} &= I - K_{t,i} Z_{t,i}^\top F_{t,i}^{-1}, \end{aligned}$$

for $t = n, \ldots, 1$ and $i = p_t, \ldots, 1$, with $r_{n,p_n} = 0$ and $N_{n,p_n} = 0$. From these recursions, we get state smoothing recursions

$$\begin{aligned} \hat{\alpha}_t &= a_{t,1} + P_{t,1} r_{t,0} \\ V_t &= P_{t,1} - P_{t,1} N_{t,0} P_{t,1}, \end{aligned}$$

and disturbance smoothing recursions

$$\begin{aligned} \hat{\epsilon}_{t,i} &= \sigma_{t,i}^2 F_{t,i}^{-1} (v_{t,i} - K_{t,i}^\top r_{t,i}), \\ \mathsf{VAR}(\hat{\epsilon}_{t,i}) &= \sigma_{t,i}^2 - \sigma_{t,i}^4 (F_{t,i}^{-1} - K_{t,i}^\top N_{t,i} K_{t,i} F_{t,i}^{-2}), \\ \hat{\eta}_t &= Q_t R_t^\top r_{t,0}, \\ \mathsf{VAR}(\hat{\eta}_{t,i}) &= Q_t R_t^\top N_{t,0} R_t Q_t. \end{aligned}$$

The recursions for diffuse phase are as follows.

$$
\begin{aligned}
L_{\infty,t,i} &= I - K_{\infty,t,i} Z_{t,i} F_{\infty,t,i}^{-1},\\
L_{t,i} &= (K_{\infty,t,i} F_{t,i} F_{\infty,t,i}^{-1} - K_{t,i}) Z_{t,i} F_{\infty,t,i}^{-1},\\
r_{0,t,i-1} &= L_{\infty,t,i}^\top r_{0,t,i},\\
r_{1,t,i-1} &= Z_{t,i}^\top v_{t,i} F_{\infty,t,i}^{-1} + L_{\infty,t,i}^\top r_{1,t,i} + L_{t,i}^\top r_{0,t,i},\\
N_{0,t,i-1} &= L_{\infty,t,i}^\top N_{0,t,i} L_{\infty,t,i}\\
N_{1,t,i-1} &= L_{t,i}^\top N_{0,t,i} L_{\infty,t,i} + L_{\infty,t,i}^\top N_{1,t,i} L_{\infty,t,i} + Z_{t,i}^\top Z_{t,i} F_{\infty,t,i}^{-1},\\
N_{2,t,i-1} &= L_{t,i}^\top N_{0,t,i} L_{t,i} + L_{\infty,t,i}^\top N_{1,t,i} L_{t,i} + (L_{\infty,t,i}^\top N_{1,t,i} L_{t,i})^\top + L_{\infty,t,i} N_{2,t,i}^\top L_{\infty,t,i}\\
&\quad - Z_{t,i}^\top Z_{t,i} F_{t,i} F_{\infty,t,i}^{-2},\\
N_{t-1,p_t} &= T_{t-1}^\top N_{t,0} T_{t-1},
\end{aligned}
$$

if $F_{\infty,t,i} > 0$, and

$$
\begin{aligned}
L_{t,i} &= I - K_{t,i} Z_{t,i} F_{t,i}^{-1},\\
r_{0,t,i-1} &= Z_{t,i}^\top v_{t,i} F_{t,i}^{-1} + L_{t,i}^\top r_{0,t,i},\\
r_{1,t,i-1} &= L_{t,i}^\top r_{1,t,i},\\
N_{0,t,i-1} &= L_{t,i}^\top N_{0,t,i} L_{t,i} + Z_{t,i}^\top Z_{t,i} F_{t,i}^{-1}\\
N_{1,t,i-1} &= N_{1,t,i} L_{t,i}\\
N_{2,t,i-1} &= N_{2,t,i} L_{t,i},
\end{aligned}
$$

otherwise. The transition from time t to $t-1$ is by $N_{j,t-1,p_t} = T_{t-1}^\top N_{j,t,0} T_{t-1}$ for $j = 0, 1, 2$, and $r_{j,t-1,p_t} = T_{t-1}^\top r_{j,t,0}$ for $j = 0, 1$, with $r_{0,d,j} = r_{d,j}$, $r_{1,d,j} = 0$, $N_{0,d,j} = N_{d,j}$, and $N_{1,d,j} = N_{2,d,j} = 0$, where (d, j) is the last point of diffuse phase. From these basic recursions, we get state smoothing recursions for diffuse phase as

$$
\begin{aligned}
\hat{\alpha}_t &= a_{t,1} + P_{t,1} r_{0,t,0} + P_{\infty,t,1} r_{1,t,0},\\
V_t &= P_{t,1} - P_{t,1} N_{0,t,0} P_{t,1} - (P_{\infty,t,1} N_{1,t,0} P_{t,1})^\top - P_{\infty,t,1} N_{1,t,0} P_{t,1} - P_{\infty,t,1} N_{2,t,0} P_{\infty,t,1},
\end{aligned}
$$

and disturbance smoothing recursions

$$
\begin{aligned}
\hat{\epsilon}_{t,i} &= -\sigma_{t,i}^2 K_{\infty,t,i}^\top r_{0,t,i},\\
\mathsf{VAR}(\hat{\epsilon}_{t,i}) &= \sigma_{t,i}^2 - \sigma_{t,i}^4 K_{\infty,t,i}^\top N_{0,t,i} K_{\infty,t,i} F_{\infty,t,i}^{-2},
\end{aligned}
$$

if $F_{\infty,t,i} > 0$, and

$$
\begin{aligned}
\hat{\epsilon}_{t,i} &= -\sigma_{t,i}^2 (v_{t,i} F_{\infty,t,i}^{-1} - K_{t,i}^\top r_{0,t,i}),\\
\mathsf{VAR}(\hat{\epsilon}_{t,i}) &= \sigma_{t,i}^2 - \sigma_{t,i}^4 (F_{t,i}^{-1} - K_{t,i}^\top N_{0,t,i} K_{t,i} F_{t,i}^{-2}),
\end{aligned}
$$

if $F_{\infty,t,i} = 0$. For $\hat{\eta}$, recursions are

$$
\begin{aligned}
\hat{\eta}_t &= Q_t R_t^\top r_{0,t,0},\\
\mathsf{VAR}(\hat{\eta}_{t,i}) &= Q_t R_t^\top N_{0,t,0} R_t Q_t.
\end{aligned}
$$

References

Anderson B, Moore J (1979). *Optimal Filtering*. Prentice-Hall, Englewood Cliffs.

Bates D, Mächler M, Bolker B, Walker S (2015a). "Fitting Linear Mixed-Effects Models Using **lme4**." *Journal of Statistical Software*, **67**(1), 1–48. doi:10.18637/jss.v067.i01.

Bates D, Maechler M (2015). ***Matrix**: Sparse and Dense Matrix Classes and Methods*. R package version 1.2-3, URL https://CRAN.R-project.org/package=Matrix.

Bates D, Maechler M, Bolker BM, Walker S (2015b). ***lme4**: Linear Mixed-Effects Models Using Eigen and S4*. R package version 1.1-10, URL http://CRAN.R-project.org/package=lme4.

Chowdhury KR (2015). ***rucm**: Implementation of Unobserved Components Model (UCM) in R*. R package version 0.6, URL http://CRAN.R-project.org/package=rucm.

Dunsmuir WTM, Scott DJ (2015). "The **glarma** Package for Observation-Driven Time Series Regression of Counts." *Journal of Statistical Software*, **67**(7), 1–36. doi:10.18637/jss.v067.i07.

Durbin J, Koopman SJ (2000). "Time Series Analysis of Non-Gaussian Observations Based on State Space Models from Both Classical and Bayesian Perspectives." *Journal of Royal Statistical Society B*, **62**, 3–56.

Durbin J, Koopman SJ (2002). "A Simple and Efficient Simulation Smoother for State Space Time Series Analysis." *Biometrika*, **89**, 603–615.

Durbin J, Koopman SJ (2012). *Time Series Analysis by State Space Methods*. 2nd edition. Oxford University Press, New York.

Eddelbuettel D (2013). *Seamless R and C++ Integration with **Rcpp***. Springer-Verlag, New York.

Eddelbuettel D, François R (2011). "**Rcpp**: Seamless R and C++ Integration." *Journal of Statistical Software*, **40**(8), 1–18. URL http://www.jstatsoft.org/v40/i08/.

Francke MK, Koopman SJ, De Vos AF (2010). "Likelihood Functions for State Space Models with Diffuse Initial Conditions." *Journal of Time Series Analysis*, **31**(6), 407–414. doi:10.1111/j.1467-9892.2010.00673.x.

Harvey AC (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.

Harvey AC, Koopman SJ (1992). "Diagnostic Checking of Unobserved-Components Time Series Models." *Journal of Business & Economic Statistics*, **10**(4), 377–89.

Holmes E, Ward E, Wills K (2013). ***MARSS**: Multivariate Autoregressive State-Space Modeling*. R package version 3.9, URL http://cran.r-project.org/web/packages/MARSS/.

Holmes EE, Ward EJ, Wills K (2012). "**MARSS**: Multivariate Autoregressive State-space models for Analyzing Time-series Data." *The R Journal*, **4**(1), 30.

King AA, Ionides EL, Bretó CM, Ellner SP, Ferrari MJ, Kendall BE, Lavine M, Nguyen D, Reuman DC, Wearing H, Wood SN (2015a). ***pomp***: *Statistical Inference for Partially Observed Markov Processes*. R package, version 1.2.1.1, URL http://kingaa.github.io/pomp.

King AA, Nguyen D, Ionides EL (2015b). "Statistical Inference for Partially Observed Markov Processes via the R Package **pomp**." *Journal of Statistical Software*, **in press**.

Koopman SJ, Durbin J (2000). "Fast Filtering and Smoothing for Multivariate State Space Models." *Journal of Time Series Analysis*, **21**, 281–296. doi:10.1111/1467-9892.00186.

Koopman SJ, Durbin J (2003). "Filtering and Smoothing of State Vector for Diffuse State-Space Models." *Journal of Time Series Analysis*, **24**, 85–98. doi:10.1111/1467-9892.00294.

Koopman SJ, Lucas A, Scharth M (2015). "Predicting Time-Varying Parameters with Parameter-Driven and Observation-Driven Models." *Review of Economics and Statistics*. doi:10.1162/REST_a_00533.

Kourentzes N, Petropoulos F (2015). ***tsintermittent***: *Intermittent Time Series Forecasting*. R package version 1.8, URL https://CRAN.R-project.org/package=tsintermittent.

Liboschik T, Fried R, Fokianos K, Probst P (2015). ***tscount***: *Analysis of Count Time Series*. R package version 1.0.0, URL https://CRAN.R-project.org/package=tscount.

Lindgren F, Rue H (2015). "Bayesian Spatial Modelling with R-**INLA**." *Journal of Statistical Software*, **63**(19), 1–25. URL http://www.jstatsoft.org/v63/i19/.

Margolin B, Kaplan N, Zeiger E (1981). "Statistical Analysis of the Ames Salmonella." *Proceedings of the National Academy of Sciences of the United States of America*, **78**(6), 3779–3783.

McCullagh P, Nelder JA (1989). *Generalized Linear Models*. 2nd edition. Chapman & Hall, London.

Petris G, Petrone S (2011). "State Space Models in R." *Journal of Statistical Software*, **41**(4), 1–25. ISSN 1548-7660. URL http://www.jstatsoft.org/v41/i04.

R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL http://www.R-project.org/.

Rigby RA, Stasinopoulos DM (2005). "Generalized Additive Models for Location, Scale and Shape." *Applied Statistics*, **54**, 507–554.

Rue H, Martino S, Lindgren F, Simpson D, Riebler A, Krainski ET (2015). ***INLA***: *Functions Which Allow to Perform Full Bayesian Analysis of Latent Gaussian Models Using Integrated Nested Laplace Approximaxion*. R package version 0.0-145253558.

Stasinopoulos M, Rigby B, Eilers P (2015). ***gamlss.util***: *GAMLSS Utilities*. R package version 4.3-2, URL https://CRAN.R-project.org/package=gamlss.util.

Statistics Finland (2014a). "Deaths by Gender, Age and Underlying Cause of Death 1969–2013." URL http://pxnet2.stat.fi/PXWeb/pxweb/en/StatFin/.

Statistics Finland (2014b). "Population According to Age (5-year) and Sex in the Whole Country 1865–2014." URL http://pxnet2.stat.fi/PXWeb/pxweb/en/StatFin/.

Szymanski C (2014). ***dlmodeler**: Generalized Dynamic Linear Modeler*. R package version 1.4-2, URL http://CRAN.R-project.org/package=dlmodeler.

Tusell F (2011). "Kalman Filtering in R." *Journal of Statistical Software*, **39**(2), 1–27. ISSN 1548-7660. URL http://www.jstatsoft.org/v39/i02.

Vasileios S (2015). ***acp**: Autoregressive Conditional Poisson*. R package version 2.1, URL https://CRAN.R-project.org/package=acp.

Yang M, Zamba GKD, Cavanaugh JE (2014). ***ZIM**: Zero-Inflated Models for Count Time Series with Excess Zeros*. R package version 1.0.2, URL https://CRAN.R-project.org/package=ZIM.

Yee TW (2010). "The **VGAM** Package for Categorical Data Analysis." *Journal of Statistical Software*, **32**(10), 1–34. URL http://www.jstatsoft.org/v32/i10/.

Affiliation:

Jouni Helske
University of Jyväskylä
Department of Mathematics and Statistics
40014 Jyväskylä, Finland
E-mail: Jouni.Helske@jyu.fi